

EXHIBIT A

ON THE LOGIC OF TLA^+

Stephan MERZ

INRIA Lorraine, LORIA, Nancy
e-mail: Stephan.Merz@loria.fr

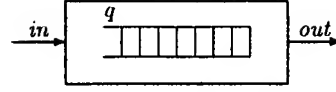
Abstract. TLA^+ is a language intended for the high-level specification of reactive, distributed, and in particular asynchronous systems. Combining the linear-time temporal logic TLA and classical set-theory, it provides an expressive specification formalism and supports assertional verification.

1 A TASTE OF TLA^+

The specification language TLA^+ has been introduced by Leslie Lamport [22] for the description of reactive and distributed, especially asynchronous systems. In this paper, I describe the semantical base of TLA^+ , which combines the linear-time temporal logic TLA and Zermelo-Fränkel set theory. My intention is not to define a new or extended formalism nor to explain the use of TLA^+ in practice. Lamport's original work covers much more material than this paper. In particular, his recent book [27] includes a tutorial introduction to writing specifications in TLA^+ , formally defines the language of TLA^+ , and describes the tools that support it. In contrast, this presentation of TLA^+ emphasizes the mathematical machinery underlying TLA^+ , explaining Lamport's choices from a logical perspective. It is my hope that it will find some use for purposes such as comparing specification formalisms or for constructing new tools to support system development in TLA^+ .

Before we begin exploring the semantics of TLA^+ , let us have a look at a simple example that introduces the typical structure of a TLA^+ specification. The TLA^+ module *SyncQueueInternal*, shown in figure 1(b), describes an unbounded FIFO queue, which is illustrated in figure 1(a). The external interface consists of an input channel *in* and an output channel *out*. Internally, the FIFO maintains a queue *q* of values that have been received via *in* but have not yet been sent via *out*.

The module consists of three sections, separated by horizontal bars for better readability, that contain declarations, definitions, and assertions. This structure of a module is conventional, but not mandatory: formally, a module is simply a list of



(a) Pictorial representation.

MODULE <i>SyncQueueInternal</i>	
EXTENDS <i>Sequences</i>	
CONSTANT <i>Message</i>	
VARIABLES <i>in, out, q</i>	
<i>NoMsg</i>	$\triangleq \text{CHOOSE } x : x \notin \text{Message}$
<i>Init</i>	$\triangleq q = \langle \rangle \wedge in = \text{NoMsg} \wedge out = \text{NoMsg}$
<i>Enq(m)</i>	$\triangleq \wedge in' \neq m$ $\wedge in' = m \wedge q' = \text{Append}(q, m)$ $\wedge out' = out$
<i>Deq</i>	$\triangleq \wedge q \neq \langle \rangle$ $\wedge out' = \text{Head}(q) \wedge q' = \text{Tail}(q)$ $\wedge in' = in$
<i>Next</i>	$\triangleq (\exists m \in \text{Message} : \text{Enq}(m)) \vee \text{Deq}$
<i>vars</i>	$\triangleq \langle in, out, q \rangle$
<i>FifoI</i>	$\triangleq \text{Init} \wedge \Box [\text{Next}]_{\text{vars}} \wedge \text{WF}_{\text{vars}}(\text{Deq})$
THEOREM	
<i>FifoI</i> $\Rightarrow \wedge \Box (q \in \text{Seq}(\text{Message}))$	
$\wedge \Box [\text{Deq} \Rightarrow out' \neq out]_{\text{vars}}$	
$\wedge \forall m \in \text{Message} : in = m \rightsquigarrow out = m$	

(b) TLA⁺ specification with the internal behavior exposed.

MODULE <i>SyncQueue</i>	
CONSTANT <i>Message</i>	
VARIABLES <i>in, out</i>	
<i>Internal(q)</i>	$\triangleq \text{INSTANCE } \text{SyncQueueInternal}$
<i>Fifo</i>	$\triangleq \exists q : \text{Internal}(q) \wedge \text{FifoI}$

(c) TLA⁺ interface specification.

Fig. 1. A FIFO queue with synchronous communication.

statements. Any identifier must have been declared or defined exactly once (possibly in an imported module) before it is used.

The first section declares *SyncQueueInternal* to be based on the standard TLA⁺ module *Sequences*, which defines finite sequences and associated operations. Next, we find a declaration of the module parameters. The constant parameter *Message* intuitively represents the set of messages that are to be sent via the FIFO queue. The variable parameters *in*, *out*, and *q* represent the current state of the queue as shown in figure 1(a); their values will change as messages are received and forwarded.

The second section contains a list of definitions, which constitute the main body of the specification. The constant *NoMsg* is defined to equal some value that is not an element of the set *Message* (section 4 explains why this definition is sensible). The state predicate *Init* identifies legal initial states of the specification: the value of *q* should be the empty sequence $\langle \rangle$, while both *in* and *out* should equal the value *NoMsg*. For any value *m*, the formula *Enq(m)* characterizes state transitions that correspond to an “enqueue” action¹: we require *m* to be different from the current value of *in* so that the queue can recognize that the input channel has changed. (This condition is not essential, but is introduced mainly for expository purposes. An implementation could for example instantiate the parameter *Message* by a set of pairs consisting of the underlying data and an extra bit, which serves to distinguish two successive enqueue actions for the same data.) The value of the variable *in* at the state following the transition, denoted by *in'*, will be *m*, and the new value of *q* is obtained by appending *m* at the end of whatever value *q* contains before the transition. Finally, we stipulate that the output channel *out* should not change during an enqueue action. The definition of the dequeue action *Deq* is similar. The action *Next* is defined as the disjunction of all enqueue actions *Enq(m)*, for *m* in *Message*, and of the dequeue action *Deq*.

The main definition of module *SyncQueueInternal* is that of the temporal formula *FifoI*, representing the “internal” specification of the FIFO queue. It is written as a conjunction: the first conjunct *Init* asserts that the first state of any behavior satisfying *FifoI* must respect the initial condition. The second conjunct specifies the next-state relation of the queue. More precisely, it asserts that every transition allowed by *FifoI* must either respect the formula *Next* or leave the expression *vars* unchanged; the latter is defined as the tuple $\langle in, out, q \rangle$ containing the state variables of interest. Because the value of a tuple is unchanged if and only if all its components are unchanged, this formula admits “stuttering steps” that do not affect the variables of interest. In a larger system that contains the FIFO queue as a component, such steps may represent actions of different components. The final conjunct of formula *FifoI* asserts a condition of weak fairness concerning the action *Deq*. It rules out behaviors where from some state onward, the *Deq* action is always enabled, but never occurs. Section 2 explores in more detail the temporal

¹ In this formula and throughout the paper, we use a standard TLA⁺ notation that displays multi-line conjunctions and disjunctions as a list “bulleted” by the connective. This layout makes long formulas easier to read and reduces the number of parentheses.

logic TLA that underlies TLA^+ , and discusses the fundamental concept of stuttering invariance.

The third section of module *SyncQueueInternal* asserts a theorem: it claims that the formula shown follows from the definitions. (In general, a module may state assumptions about constants, and theorems should then follow from the definitions and the stated or imported assumptions.) In a loose reading, the assertion of a theorem in a module can be regarded merely as a comment that highlights the specifier's intuitions. Formally, however, a theorem represents a proof obligation that must be discharged for the module to be correct, and we will turn to proof rules for verification in section 3. The theorem asserted of module *SyncQueueInternal* states that every behavior that satisfies formula *FifoI* has the following properties:

- at every state, the value of the variable q is a finite sequence whose elements are contained in *Message*,
- every *Deq* step changes the value of the output channel *out*, and
- every message that appears on the input channel will eventually be output.

The current version of TLA^+ described in [27] does not contain a language for writing proofs, although Lamport advocates a hierarchical proof notation [26].

Like HOL [16] and other logical specification languages, TLA^+ is declarative: the names of formulas such as *Init*, *Next* or *FifoI* are formally irrelevant, although it is good practice to make them meaningful. The meaning of a formula can always be uniquely and compositionally determined from the meaning of its subformulas, and how to do this is the main subject of the present paper. As in any logic, there are many logically equivalent ways to express a specification. For example, we could have replaced the definitions of *Enq* and *Next* by

$$\begin{aligned} \text{Enq} &\triangleq \wedge in' \in \text{Message} \wedge in' \neq in \\ &\quad \wedge q' = \text{Append}(q, in') \wedge out' = out \\ \text{Next} &\triangleq \text{Enq} \vee \text{Deq} \end{aligned}$$

without changing the meaning of formula *FifoI*.

TLA^+ does not hide the complexity of a system by using built-in data types; as we will see in section 4, every value is just a set. Similarly, it does not presuppose any fixed system model such as shared-variable or message-passing concurrency, synchronous or asynchronous communication, etc. Its expressiveness comes from the unfettered use of set theory and the mechanism of definitional extension. For our example, we have chosen the internal variable q to change at the same time as the interface variables *in* and *out*, representing a synchronous style of communication. A specification of a FIFO queue using asynchronous communication channels is presented in Lamport's book [27, ch. 4].

The specifications of module *SyncQueueInternal* describe the behavior of the FIFO queue in terms of the three variables *in*, *out*, and q . One important principle in writing specifications is that of *information hiding*, which requires a component

specification not to reveal the internal structure (or “implementation details”) of the component. In our example, the variable q is such an implementation detail: as illustrated by the box in figure 1(a), only the behavior of the externally visible variables in and out should be constrained by the queue specification. Module *SyncQueue*, shown in figure 1(c), contains an interface specification of the FIFO queue based on the previous specification. In fact, it declares in and out as its only variable parameters. The following line instantiates the previously discussed module *SyncQueueInternal*: any operator Op defined in that module can be referenced as *Internal*(q)! Op in module *SyncQueue*. The general form of instantiation in TLA⁺ allows for substitution of expressions for module parameters; any remaining parameters are implicitly instantiated with the identifier of the same name valid at the point of instantiation; it is an error if that identifier has not been declared or defined. In our case, the parameters *Message*, in , and out of module *SyncQueueInternal* are instantiated by the corresponding parameters of module *SyncQueue*, whereas parameter q is instantiated by the local parameter of the operator *Internal*.

Module *SyncQueue* then defines the formula *Fifo*, representing the interface specification of the FIFO queue, as the formula obtained from *Internal*(q)!*FifoI* by existential quantification over q . This formula is satisfied by every behavior where in and out take the values as described by the internal specification, but where q may take arbitrary values. (The precise semantics is defined in section 2.4.) In this respect, existential quantification represents hiding of internal state components, and formula *Fifo* specifies the interface of the FIFO queue.

2 TLA: THE TEMPORAL LOGIC OF ACTIONS

TLA⁺ combines TLA, the Temporal Logic of Actions [25], and mathematical set theory. We now present the semantics of TLA, while sections 3 and 4 explore the verification of temporal formulas and the specification of data structures in set theory. Again, we emphasize that this exposition is aimed at a precise definition of TLA as a logical language; it does not attempt to explain how TLA is used to specify algorithms or systems.

2.1 Rationale

The logic of time has its origins in philosophy and linguistics, where it was intended to formalize temporal references in natural language [21, 34]. Around 1975, Pnueli [33] and others recognized that such logics could be useful as a basis for the semantics of computer programs. In particular, traditional formalisms based on pre- and post-conditions were found to be ill-suited for the description of reactive systems that are continuously interacting with their environment and that are not necessarily intended to terminate. Temporal logic, as it came to be called in computer science, offered an elegant framework to describe safety and liveness properties [10, 24] of reactive systems. Different dialects of temporal logic can be

distinguished according to the properties assumed of the underlying model of time (e.g., discrete or dense) and to the connectives offered to refer to different moments in time (e.g., future vs. past references). For computer science applications, the most controversial distinction has been between linear-time and branching-time logics. In the linear-time view, a system is identified with the set of its executions, modeled as infinite sequences of states, whereas the branching-time view also considers the branching structure of a system. Linear-time temporal logics, including TLA, suffice to formulate correctness properties that hold of all the runs of a system, whereas branching-time temporal logics can also express possibility properties such as the existence of a path, from every reachable state, to a “reset” state. The discussion of the relative merits and deficiencies of these two kinds of temporal logics is beyond the scope of this paper, but see, e.g., Vardi [38] for a recent contribution to this subject, with many references to earlier papers.

Despite initial enthusiasm about the usefulness of temporal logic as a language to describe individual system properties, attempts to actually write complete system specifications in temporal logic revealed that not even a component as simple as a FIFO queue could be unambiguously specified [35]. This observation has led many researchers to propose that reactive systems should be modeled by some variant of state machines while temporal logic was retained as a high-level language to describe the correctness properties. A major breakthrough came with the insight that temporal logic properties are decidable for finite-state models, and such model checking techniques [13] are nowadays routine for the debugging of hardware circuits and communication protocols.

Another weakness of standard temporal logic becomes apparent when one attempts to compare two specifications of the same system, written at different levels of abstraction. Specifically, atomic system actions are usually described via a “next-state” operator, but the “grain of atomicity” typically changes during refinement, complicating comparisons between specifications. For example, we might want to refine the specification of the FIFO queue of figure 1(b) such that the operation of appending an element to a queue is described as a sequence of more elementary assignments.

TLA has been designed as a formalism where system specifications and their properties are expressed in the same language, and where refinement is reduced to elementary logic. The problems mentioned above are addressed in the following ways: “internal” specifications are written by defining their initial conditions and next-state relations, resembling the description of state machines, and are augmented by liveness and fairness conditions. Abstractness in the sense of information hiding is ensured by quantification over state variables. The refinement problem is solved by systematically allowing for stuttering steps that do not change the values of the state variables of interest; an implementation is allowed to refine such high-level stuttering into lower-level state changes. Similar ideas can be found in Back’s refinement calculus [11] and in more recent versions of Abrial’s B method [9]. However, in order to prevent infinite stuttering, these formalisms require side conditions that are expressed in terms of well-founded orderings. Temporal logic can

state such requirements more abstractly in terms of high-level fairness conditions that must be preserved by a refinement, using any combination of fairness conditions and arguments based on well-founded orderings.

Based on these concepts, TLA provides a unified logical language to express system specifications and their properties. A single set of logical rules is used for system verification and for proving refinement.

2.2 Transition formulas

The language of TLA is two-tiered: the base tier contains formulas that describe states and state transitions, whereas the top tier consists of temporal formulas that are evaluated over infinite sequences of states. In this section, we define the syntax and semantics of transition formulas, whereas the following sections will consider temporal formulas. Because transition formulas are just ordinary (untyped, first-order) predicate logic, this section can be quite brief.

Assume a given signature of first-order predicate logic, consisting of:

- at most denumerable sets \mathcal{L}_F and \mathcal{L}_P of function and predicate symbols, each symbol equipped with its arity, and
- a denumerable set \mathcal{V} of variables, partitioned into denumerable sets \mathcal{V}_F and \mathcal{V}_R of flexible and rigid variables.

These sets should be disjoint from one another; moreover, no variable in \mathcal{V} should be of the form v' . By $\mathcal{V}_{F'}$, we denote the set $\{v' \mid v \in \mathcal{V}_F\}$ of primed flexible variables, and by \mathcal{V}_E , the union $\mathcal{V} \cup \mathcal{V}_{F'}$ of primed and unprimed variable symbols.

Transition functions and *transition predicates* (also called *actions*) are terms and formulas built from the symbols in \mathcal{L}_F and \mathcal{L}_P , and from the variables in \mathcal{V}_E . For example, if f is a ternary function symbol, p is a unary predicate symbol, $x \in \mathcal{V}_R$, and $v \in \mathcal{V}_F$, then the term e defined as $f(v, x, v')$ is a transition function, and the formula C defined as $\exists v' : p(f(v, x, v')) \wedge \neg(v' = x)$ is an action. We omit a formal inductive definition of the syntax of transition functions and formulas. Collectively, we use the term *transition formula* to refer to transition functions and predicates.

The semantics of transition formulas is also unsurprising. It is based on a first-order interpretation, which defines a universe of values and interprets each symbol in \mathcal{L}_F by a function and each symbol in \mathcal{L}_P by a relation of appropriate arities. In preparation for the semantics of temporal formulas, we distinguish between the valuations of flexible and rigid variables. A *state* is a mapping of the flexible variables in \mathcal{V}_F to values of the universe. Given two states s and t and a valuation ξ of the rigid variables in \mathcal{V}_R , we can define the valuation $\alpha_{s,t,\xi}$ of the variables in \mathcal{V}_E as the mapping such that $\alpha_{s,t,\xi}(x) = \xi(x)$ for $x \in \mathcal{V}_R$, $\alpha_{s,t,\xi}(v) = s(v)$ for $v \in \mathcal{V}_F$, and $\alpha_{s,t,\xi}(v') = t(v)$ for $v' \in \mathcal{V}_{F'}$. The semantics of a transition function or transition formula E , written $\llbracket E \rrbracket_{s,t}^\xi$, is then simply the standard predicate logic semantics of E with respect to the extended valuation $\alpha_{s,t,\xi}$.

We say that a transition predicate A is *valid* for the interpretation iff $\llbracket A \rrbracket_{s,t}^\xi$ is true for all states s, t and all valuations ξ . It is *satisfiable* iff $\llbracket A \rrbracket_{s,t}^\xi$ is true for some s, t , and ξ . Similarly, A is valid (satisfiable) for a class \mathcal{C} of interpretations iff it is valid for all (satisfiable for some) interpretations in \mathcal{C} .

Finally, the notions of free and bound variables in a transition formula are defined as usual, as is the notion of substitution of a transition function a for a variable v , written $E[a/v]$. We assume that capture of free variables in a substitution is avoided by an implicit renaming of bound variables. For example, the set of free variables of the transition function e shown above is $\{x, v, v'\}$, and v' is a bound variable of the action C . We emphasize that at the level of transition formulas, we consider v and v' to be distinct, unrelated variables.

State formulas are transition formulas that do not contain free primed flexible variables. For example, the action C above is actually a state predicate. Because the semantics of state formulas only depends on a single state, we simply write $\llbracket P \rrbracket_s^\xi$ when P is a state formula. The subclass of *constant formulas* is even more restrictive in that only free occurrences of rigid variables are allowed; consequently, their semantics depends only on the valuation ξ . (Arguably, *rigid formulas* would be a more appropriate name for this class, but the TLA literature consistently uses the designation constant formulas.)

TLA introduces some specific abbreviations at the level of transition formulas. If E is a state formula then E' is the transition formula obtained from E by replacing each free occurrence of a flexible variable v in E with its primed counterpart v' (where bound variables are renamed as necessary). For example, since C is a state formula, we may build the formula C' by substituting v' for v . Since v' is bound in C , this results in the formula $\exists y : p(f(v', x, y)) \wedge \neg(y = x)$, up to renaming of the bound variable.

For an action A , the state formula $\text{ENABLED } A$ is obtained by existential quantification over all primed flexible variables that occur free in A . Thus, $\llbracket \text{ENABLED } A \rrbracket_s^\xi$ holds if $\llbracket A \rrbracket_{s,t}^\xi$ holds for some state t , that is, if action A may occur in state s . For actions A and B , the action $A \cdot B$ is defined as $\exists z : A[z/v'] \wedge B[z/v]$ where v is a list of all flexible variables v_i such that v_i occurs free in B or v'_i occurs free in A , and z is a corresponding list of fresh variables. It follows that $\llbracket A \cdot B \rrbracket_{s,t}^\xi$ holds iff both $\llbracket A \rrbracket_{s,u}^\xi$ and $\llbracket B \rrbracket_{u,t}^\xi$ hold for some state u .

Because these abbreviations are defined in terms of quantification and substitution, their interplay can be quite delicate. For example, $\text{ENABLED } P$ is by definition just P for any state predicate P , and therefore $(\text{ENABLED } P)'$ equals P' . On the other hand, $\text{ENABLED } (P')$ is a constant formula—if P does not contain any rigid variables then $\text{ENABLED } (P')$ is valid iff P is satisfiable.

For an action A and a state function t we write $[A]_t$ to denote $A \vee t' = t$, and $\langle A \rangle_t$ for $A \wedge \neg(t' = t)$. Therefore, $[A]_t$ requires A to hold only if t changes value during a transition, whereas the dual formula $\langle A \rangle_t$ strengthens A in requiring that t changes value while A holds true.

2.3 Temporal formulas

We now turn to the temporal tier of TLA. Because it is less familiar than first-order predicate logic and because we wish to give precise definitions, we devote much more space to its presentation. However, the temporal formulas that one actually writes in TLA⁺ specifications usually follow a standard idiom, and more than 95% of a typical specification consist of definitions at the transition level.

The (temporal) formulas of TLA are inductively defined as follows:

- Every state formula is a formula.
- Boolean combinations (connectives including \neg , \wedge , \vee , \Rightarrow , and \equiv) of formulas are formulas.
- If F is a formula then so is $\Box F$ (“always F ”).
- If A is an action and t is a state function then $\Box[A]_t$ (“always square A sub t ”) is a formula.
- If F is a formula and x is a rigid variable then $\exists x : F$ is a formula.
- If F is a formula and v is a flexible variable then $\exists v : F$ is a formula.

In particular, an action A by itself is not a temporal formula, not even in the form $[A]_t$. Actions can occur only in subformulas $\Box[A]_t$.

To determine free and bound variables at the temporal level, we do not distinguish between primed and unprimed occurrences of flexible variables, and the quantifier \exists binds both kinds of occurrences. More formally, the set of free variables of a temporal formula is a subset of $\mathcal{V}_F \cup \mathcal{V}_R$. The free occurrences of variables in a state formula P , considered as a temporal formula, are precisely the free occurrences in P , considered as a transition formula. However, variable $v \in \mathcal{V}_F$ has a free occurrence in $\Box[A]_t$ iff either v or v' has a free occurrence in A or in t . Similarly, substitution $F[e/v]$ of a state function e for a flexible variable v substitutes both e for v and e' for v' in the action subformulas of F , again after renaming bound variables as necessary. For example, substitution of the state function $h(v)$, where $h \in \mathcal{L}_F$ and $v \in \mathcal{V}_F$, for w in the temporal formula

$$\exists v : p(v, w) \wedge \Box[q(v, f(w, v'), w')]_{g(v, w)}$$

results in the formula (up to renaming of the bound variable)

$$\exists u : p(u, h(v)) \wedge \Box[q(u, f(h(v), u'), h(v'))]_{g(u, h(v))}$$

Because state formulas do not contain free occurrences of primed flexible variables, the definitions of substitutions for transition formulas and for temporal formulas agree on state formulas. The substitution of a (proper) transition function for a variable is not allowed as it could result in an expression that is not a well-formed TLA formula.

The semantics of temporal formulas is defined in terms of *behaviors*, which are infinite sequences of states, and of valuations of the rigid variables. For a behavior $\sigma = s_0 s_1 \dots$, we write σ_i to refer to state s_i , and $\sigma|_i$ to denote the suffix $s_i s_{i+1} \dots$. The following inductive definition assigns a truth value $\llbracket F \rrbracket_\sigma^\xi \in \{t, f\}$ to every formula F :

- $\llbracket P \rrbracket_\sigma^\xi = \llbracket P \rrbracket_{\sigma_0}^\xi$: state formulas are evaluated at the initial state of the behavior.
- The semantics of Boolean operators is the usual one.
- $\llbracket \Box F \rrbracket_\sigma^\xi = t$ iff $\llbracket F \rrbracket_{\sigma|_i}^\xi = t$ for all $i \in \mathbb{N}$: this is the usual clause from linear-time temporal logic.
- $\llbracket \Box[A]_t \rrbracket_\sigma^\xi = t$ iff for all $i \in \mathbb{N}$, $\llbracket t \rrbracket_{\sigma|_i}^\xi = \llbracket t \rrbracket_{\sigma|_{i+1}}^\xi$ or $\llbracket A \rrbracket_{\sigma|_i, \sigma|_{i+1}}^\xi = t$: such a formula holds iff every state transition in σ that does not leave t unchanged satisfies A .
- $\llbracket \exists x : F \rrbracket_\sigma^\xi = t$ iff $\llbracket F \rrbracket_\sigma^\eta = t$ for some valuation η such that $\eta(y) = \xi(y)$ for all $y \in \mathcal{V}_R \setminus \{x\}$: this is again the usual definition from predicate logic.
- The semantics of formulas $\exists v : F$ will be defined in section 2.4 below.

Abbreviations for temporal formulas include the universal quantifiers \forall and \forall over rigid and flexible variables. The formula $\Diamond F$ (“eventually F ”), defined as $\neg \Box \neg F$, asserts that F holds of some suffix of the behavior; similarly, $\Diamond[A]_t$ (“eventually angle A sub t ”) is defined as $\neg \Box[\neg A]_t$ and asserts that some future transition satisfies A and changes the value of t . We write $F \rightsquigarrow G$ (“ F leads to G ”) for the formula $\Box(F \Rightarrow \Diamond G)$, which asserts that every occurrence of F will eventually be followed by an occurrence of G . Combinations of the “always” and “eventually” operators express “infinitely often” ($\Box \Diamond$) and “almost always” ($\Diamond \Box$). Observe that a formula can be both infinitely often true and infinitely often false, thus “almost always” is strictly stronger than “infinitely often”. These combinations are especially important as the building blocks to formulate fairness conditions. In particular, weak and strong fairness for an action $\langle A \rangle_t$ are defined as

$$\begin{aligned} \text{WF}_t(A) &\triangleq (\Box \Diamond \neg \text{ENABLED } \langle A \rangle_t) \vee \Box \Diamond \langle A \rangle_t & (\equiv \Diamond \Box \text{ENABLED } \langle A \rangle_t \Rightarrow \Box \Diamond \langle A \rangle_t) \\ \text{SF}_t(A) &\triangleq (\Diamond \Box \neg \text{ENABLED } \langle A \rangle_t) \vee \Box \Diamond \langle A \rangle_t & (\equiv \Box \Diamond \text{ENABLED } \langle A \rangle_t \Rightarrow \Box \Diamond \langle A \rangle_t) \end{aligned}$$

Weak fairness stipulates that an action $\langle A \rangle_t$ occurs infinitely often during a behavior if it is almost always enabled; strong fairness even requires that the action must happen infinitely often if it is infinitely often, but not necessarily persistently, enabled.

2.4 Stuttering invariance and quantification

Formulas $\Box[A]_t$ allow for “stuttering”: besides state transitions that satisfy A , they also admit any transitions that do not change the state function t . In particular, duplications of states can not be observed by formulas of this form. Stuttering invariance is important in connection with refinement and composition [24].

To formalize this notion, for a set V of flexible variables we define two states s and t to be V -equivalent, written $s =_V t$, iff $s(v) = t(v)$ for all $v \in V$. We define V -stuttering equivalence, written \approx_V , as the smallest equivalence relation on behaviors that contains $\rho \circ \langle s \rangle \circ \sigma$ and $\rho \circ \langle t, u \rangle \circ \sigma$, for any finite sequence of states ρ , infinite sequence of states σ , and V -equivalent states $s =_V t =_V u$. Intuitively, V -stuttering equivalence allows for duplication and deletion of finite repetitions of V -equivalent states. In particular, the relation \approx_{V_F} , which we also write as \approx , identifies two behaviors that differ by duplications or deletions of identical states.

The fundamental theorem asserting that TLA is not expressive enough to distinguish stuttering-equivalent behaviors can now be formally stated as follows:

Theorem 1 (stuttering invariance). Assume that F is a TLA formula whose free flexible variables are among V , that $\sigma \approx_V \tau$ are V -stuttering equivalent behaviors, and that ξ is a valuation. Then $\llbracket F \rrbracket_\sigma^\xi = \llbracket F \rrbracket_\tau^\xi$.

For TLA formulas without quantification over flexible variables, it is not hard to prove theorem 1 from the semantic clauses of section 2.3 by induction on the structure of formulas [25, 6]. On the other hand, quantification over flexible variables requires some attention: the “obvious” semantic clause for formulas $\exists v : F$ would read $\llbracket \exists v : F \rrbracket_\sigma^\xi = t$ iff $\llbracket F \rrbracket_\tau^\xi = t$ for some behavior τ whose states τ_i agree with the corresponding states σ_i on all variables except for v . This definition, however, would not preserve stuttering invariance. For example, consider the formula

$$F \triangleq v = c \wedge w = c \wedge \Diamond(w \neq c) \wedge \Box[v \neq c]_w$$

σ

v
 w

c
 c

d
 c

d
 d

\dots
 \dots

that requires that both variables v and w initially equal the constant c , that eventually, w is different from c , and that v must be different from c whenever w changes value. Any behavior σ that satisfies F must therefore contain two distinct transitions, the first of which changes v from c to some other value (but preserving the value of w), while the second transition changes w , as indicated in the picture. In particular, $\sigma_1(w)$ must equal c , hence the above definition of quantification implies that $\tau_1(w)$ must equal c , for any behavior τ satisfying the formula $\exists v : F$. However, the behavior τ obtained from σ by removing the second state (where v equals d but w equals c) is $\{w\}$ -stuttering equivalent to σ . Because w is the only free flexible variable of $\exists v : F$, theorem 1 asserts that τ should satisfy $\exists v : F$, although $\tau_1(w)$ is different from c .

In other words, the definition of quantification over flexible variables must allow for the removal of transitions that modify only the bound variables. This observation motivates the following semantic clause for quantified formulas: for a flexible variable v , we say that two behaviors σ and τ are *equal up to v* iff σ_i and τ_i agree on all variables in $V_F \setminus \{v\}$, for all $i \in \mathbb{N}$. We say that σ and τ are *similar up to v* , written $\sigma \simeq_v \tau$ iff there exist behaviors σ' and τ' such that

- σ and σ' are stuttering equivalent ($\sigma \approx \sigma'$),

- σ' and τ' are equal up to v , and
- τ and τ' are again stuttering equivalent ($\tau \approx \tau'$).

Now, we define $\llbracket \exists v : F \rrbracket_\sigma^x = t$ iff $\llbracket F \rrbracket_\tau^x = t$ holds for some behavior τ similar to σ up to v . Because the definition of \simeq_v explicitly allows for stuttering, theorem 1 is preserved for all TLA formulas.

2.5 Properties, refinement, and composition

As we have seen in the example of the FIFO queue, TLA uses the same formalism of temporal logic to represent system specifications and properties. System specifications are usually written in the form

$$\exists x : \text{Init} \wedge \Box[\text{Next}]_v \wedge L$$

where v is the list of all relevant state variables, x is the list of internal (hidden) variables, Init is a state predicate representing the initial condition, Next is an action that describes the next-state relation, usually written as a disjunction of more elementary actions, and L is a conjunction of formulas $\text{WF}_v(A)$ or $\text{SF}_v(A)$ asserting fairness assumptions of individual disjuncts of Next . However, other forms of specifications are possible and can occasionally be useful. Asserting that a property F holds of a specification S amounts to saying that every behavior that satisfies S must also satisfy F ; in other words, it asserts the validity of the implication $S \Rightarrow F$. For example, the theorem asserted in module *SyncQueueInternal* states three essential properties of the FIFO queue.

Unlike most other temporal logics, TLA is intended to support stepwise system development by refinement of specifications [11]. The basic idea of refinement consists in successively adding implementation detail while preserving the properties required at an abstract level. In a refinement-based approach to system development, one proceeds by writing successive models, each of which introduces some additional detail while preserving the essential properties of the preceding model. Fundamental properties of a system can thus be established at high levels of abstraction, errors can be detected in early phases, and the complexity of formal assurance is spread over the entire development process. A refinement C preserves all TLA properties of an abstract specification A if and only if for every formula F , if $A \Rightarrow F$ is valid, then so is $C \Rightarrow F$. This condition is in turn equivalent to requiring the validity of $C \Rightarrow A$. Because C will contain extra variables to represent the lower-level detail, and because these variables will change in transitions that have no counterpart at the abstract level, stuttering invariance of TLA formulas is essential to make validity of implication a reasonable definition of refinement.

Stuttering invariance is also essential for composition to be representable as conjunction [18]. In fact, if A and B are TLA specifications of two components, then $A \wedge B$ describes those behaviors that satisfy both components' initial conditions, that allow actions of either process to occur, synchronizing on common variables

(which represent interfaces between the components), and that satisfy all relevant liveness properties. In particular, stuttering invariance ensures that each component may perform local actions without interfering with the specification of the other component.

As a test of these ideas, we might try to prove that two FIFO queues in a row again implement a FIFO queue. Let us assume that the two queues are connected by a channel mid , then the above principles seem to imply that the formula²

$$Fifo[mid/out] \wedge Fifo[mid/in] \Rightarrow Fifo$$

is valid. Unfortunately, this is not true, for the following reason: formula $Fifo$ implies that the in and out channels never change simultaneously, whereas the conjunction on the left-hand side allows such changes (if the left-hand queue performs an Enq action, while the right-hand queue performs a Deq). This technical problem can be attributed to a design decision taken in the specification of the FIFO queue to disallow simultaneous changes to its input and output interfaces, a specification style known as “interleaving specifications”. In fact, the argument merely shows that the composition of two interleaving queues does not implement a interleaving queue. Choosing an interleaving or a non-interleaving style is an artifact of the model that represents the actual system; interleaving specifications are usually easier to write and to understand. The problem disappears if we explicitly add an “interleaving” assumption for the composition: the implication

$$Fifo[mid/out] \wedge Fifo[mid/in] \wedge \Box[in' = in \vee out' = out]_{in,out} \Rightarrow Fifo \quad (1)$$

is valid and its proof will be considered in section 3.5.

2.6 Variations and extensions

We discuss some of the choices that we have made in the presentation of TLA , as well as possible extensions.

Transition formulas and priming. Our presentation of TLA is based on standard first-order logic, to the extent possible. In particular, we have defined transition formulas as formulas of ordinary predicate logic over a large set \mathcal{V}_E of variables where v and v' are unrelated. An alternative presentation would consider $'$ as an operator, resembling the next-time modality of temporal logic. In fact, this appears to be the presentation preferred by Lamport [27]. The semantics of temporal formulas is unaffected by the choice of presentation, and the style adopted in this paper corresponds well to the verification rules of TLA , explored in section 3.

² TLA^+ introduces concrete syntax, based on module instantiation, for writing substitutions such as $Fifo[mid/out]$.

Compositional verification. We have argued in section 2.5 that composition is represented in TLA as conjunction. Because components can rarely be expected to operate correctly in arbitrary environments, their specifications usually include some assumptions about the environment. An *open system specification* is one that does not constrain its environment; it asserts that the component will function correctly provided that the environment behaves as expected. One way to write such specifications is in the form of implications $E \Rightarrow M$ where E describes the environment assumptions and M , the component specification. However, it turns out that often a stronger form of specifications is desirable that requires the component to adhere to its description M for at least as long as the environment has not broken its obligation E . In particular, when systems are built from “open” component specifications, this form, written $E \dot{\Rightarrow} M$, allows for a strong composition rule that can discharge mutual assumptions between components [4, 14]. It can be shown that the operator $\dot{\Rightarrow}$ is actually definable in TLA, and that the resulting composition rule can be justified in terms of an abstract logic of specifications, supplemented by principles specific to TLA [5, 7].

TLA*. TLA defines distinct tiers of transition formulas and temporal formulas, where transition formulas must be guarded by “brackets” to ensure stuttering invariance. Although the separation between the two tiers is natural when writing system specifications, it is not a prerequisite to obtaining stuttering invariance. In [32], I have defined the logic TLA* whose syntax distinguishes the two classes of *pure* and *impure* formulas. Whereas pure formulas of TLA* contain impure formulas in the same way that temporal formulas of TLA contain transition formulas, impure formulas generalize transition formulas in that they admit Boolean combinations of F and $\circ G$, where F and G are pure formulas and \circ is the next-time modality of temporal logic. For example, the TLA* formula

$$\Box[A \Rightarrow \circ \Diamond \langle B \rangle_u]_t$$

requires that every $\langle A \rangle_t$ action must eventually be followed by $\langle B \rangle_u$. Assuming appropriate syntactic conventions, TLA* is a generalization of TLA because every TLA formula is also a TLA* formula, with the same semantics. On the other hand, it can be shown that every TLA* formula can be expressed in TLA using some additional quantifiers. For example, the TLA* formula above is equivalent to the TLA formula³

$$\begin{aligned} \exists v : & \wedge \Box((v = c) \equiv \Diamond \langle B \rangle_u) \\ & \wedge \Box[A \Rightarrow v' = c]_t \end{aligned}$$

where c is a constant and v is a fresh flexible variable. TLA* thus offers a richer syntax without increasing the expressiveness, allowing high-level requirement spec-

³ Strictly, this equivalence is true only for universes that contain at least two distinct values; one-element universes are not very interesting.

ifications to be expressed more directly. On the other hand, the propositional fragment of TLA^+ admits a rather straightforward axiomatization. (No proof system is known for propositional TLA, although Abadi [1] proposed a rather involved axiomatization of an early version of TLA that was not invariant under stuttering.) For example,

$$\Box[F \Rightarrow \circ F]_v \Rightarrow (F \Rightarrow \Box F)$$

where F is a temporal formula and v is a tuple containing all flexible variables with free occurrences in F , is a TLA^+ formulation of the usual induction axiom of temporal logic; this is a TLA formula only if F is in fact a state formula.

Binary temporal operators. Unlike standard linear-time temporal logics [30], TLA does not include binary operators such as **until**, because they are not necessary for writing system specifications, and because they can be confusing, especially when nested. These operators are, however, definable in TLA using quantification over flexible variables. For example, suppose that P and Q are state predicates whose free variables are among w , that v is a flexible variable that does not appear in w , and that c is a constant. Then P **until** Q can be defined as the formula

$$\begin{aligned} \exists v : & \wedge (v = c) \equiv Q \\ & \wedge \Box[(v \neq c \Rightarrow P) \wedge (v' = c \equiv (v = c \vee Q'))]_{(v,w)} \\ & \wedge \Diamond Q \end{aligned}$$

The idea is to use the auxiliary variable v to remember whether Q has already been true. As long as Q has been false, P is required to hold. For arbitrary TLA formulas F and G , the formula F **until** G can be defined along the same lines, using a technique as shown for the translation of TLA^+ formulas above.

3 TLA PROOF RULES

Since TLA formulas are used to describe systems as well as their properties, deductive system verification can be based on logical axioms and rules of TLA. More precisely, a system described by formula *Spec* has property *Prop* if and only if every behavior that satisfies *Spec* also satisfies *Prop*, that is, iff the implication $Spec \Rightarrow Prop$ is valid over the class of interpretations where the function and predicate symbols have the intended meaning. System verification, in principle, therefore requires reasoning about sets of behaviors. The TLA proof rules are designed to reduce this temporal reasoning, as far as possible, to the proof of verification conditions expressed in the underlying predicate logic, a strategy that is commonly referred to as *assertional reasoning*. In this section, we state some typical rules and illustrate their use; more information can be found elsewhere [25].

3.1 Invariants

Invariants characterize the set of states that can be reached during system execution; they constitute the basic safety properties of interest and are also the starting point for almost any verification attempt. In TLA, an invariant is expressed by a formula of the form $\Box I$ where I is a state formula.

A basic rule for proving invariants is given by

$$\frac{I \wedge [N]_t \Rightarrow I'}{I \wedge \Box[N]_t \Rightarrow \Box I} \text{ (INV1)}$$

This rule asserts that for every interpretation for which the antecedent $I \wedge [N]_t \Rightarrow I'$ is a valid transition formula, the consequent $I \wedge \Box[N]_t \Rightarrow \Box I$ is a valid temporal formula. The antecedent states that every possible transition (stuttering or not) preserves I ; thus, if I holds initially it is guaranteed to hold forever. Formally, the correctness of rule (INV1) is easily established by induction on behaviors. Because the antecedent is a transition formula, its proof relies on standard axioms and proof rules of predicate logic, augmented by “data” axioms that characterize the intended interpretations.

For example, we can use (INV1) to prove the invariant $\Box(q \in \text{Seq}(\text{Message}))$ of the FIFO queue specified in module *SyncQueueInternal* of figure 1(b). We have to prove

$$\text{FifoI} \Rightarrow \Box(q \in \text{Seq}(\text{Message})) \quad (2)$$

which, by rule (INV1), the definition of formula *FifoI*, and propositional logic can be reduced to proving

$$\text{Init} \Rightarrow q \in \text{Seq}(\text{Message}) \quad (3)$$

$$q \in \text{Seq}(\text{Message}) \wedge [Next]_{\text{vars}} \Rightarrow q' \in \text{Seq}(\text{Message}) \quad (4)$$

Because the empty sequence is certainly a finite sequence of messages, (3) follows from the definition of *Init* and appropriate data axioms. Similarly, the proof of (4) reduces to proving preservation of the invariant under stuttering, *Deq*, and *Enq(m)* actions, for any $m \in \text{Message}$, all of which are again straightforward.

3.2 Step simulation

The following rule can be used to prove “action invariants”; it relies on a previously proven state invariant I :

$$\frac{I \wedge I' \wedge [M]_t \Rightarrow [N]_u}{\Box I \wedge \Box[M]_t \Rightarrow \Box[N]_u} \text{ (TLA2)}$$

In particular, it follows from (TLA2) that the next-state relation can be strengthened by an invariant:

$$\Box I \wedge \Box [M]_t \Rightarrow \Box [M \wedge I \wedge I']_t$$

Note that the converse of this implication is not valid: the right-hand side holds of any behavior where t never changes, independently of the value of I .

We may use (TLA2) to prove that the FIFO queue never dequeues the same value twice in a row:

$$FifoI \Rightarrow \Box [Deq \Rightarrow out' \neq out]_{vars} \quad (5)$$

This proof requires an invariant that in particular asserts that no consecutive elements of the internal queue are identical:

$$\begin{aligned} Inv &\triangleq \text{LET } oq \triangleq (out) \circ q \\ &\quad \text{IN } \wedge in = oq[Len(oq)] \\ &\quad \wedge \forall i \in 1..Len(oq) - 1 : oq[i] \neq oq[i + 1] \end{aligned}$$

We have used some TLA^+ syntax in formulating Inv : the local abbreviation oq denotes the sequence obtained by prefixing the current value of the output channel out to the internal queue q ; also, TLA^+ represents a sequence s as a function such that its elements can be accessed as $s[1], \dots, s[Len(s)]$. Formula Inv asserts that the current value of the input channel equals the last element of the sequence oq , and that no two consecutive elements of oq are identical. The proof of $FifoI \Rightarrow \Box Inv$ follows the pattern used in proving invariant (2) above, using rule (INV1).

For the proof of (5), rule (TLA2) requires that we show the validity of

$$Inv \wedge Inv' \wedge [Next]_{vars} \Rightarrow [Deq \Rightarrow out' \neq out]_{vars} \quad (6)$$

The proof of (6) reduces to the three cases of a stuttering transition, an $Enq(m)$ action, and a Deq action. Only the last case is non-trivial. Its proof relies on the definition of Deq , which implies that q is non-empty and that $out' = Head(q)$. In particular, the sequence oq contains at least two elements, and therefore Inv implies that $oq[1]$, which is just out , is different from $oq[2]$, which is $Head(q)$. This suffices to prove $out' \neq out$.

3.3 Liveness properties

Liveness properties, intuitively, assert that something good must eventually happen [10, 23]. Because formulas $\Box[N]_t$ are satisfied by stuttering behavior and do not require any progress, the proof of liveness properties must ultimately rely on fairness properties assumed of the specification. TLA provides rules to deduce elementary liveness properties from the fairness properties assumed of a specification; more complex properties can then be inferred with the help of well-founded orderings.

The following rule can be used to prove a leads-to formula from a weak fairness assumption; a similar rule exists for strong fairness.

$$\frac{\begin{array}{l} I \wedge I' \wedge P \wedge [N]_t \Rightarrow P' \vee Q' \\ I \wedge I' \wedge P \wedge \langle N \wedge A \rangle_t \Rightarrow Q' \\ I \wedge P \Rightarrow \text{ENABLED } \langle A \rangle_t \end{array}}{\Box I \wedge \Box [N]_t \wedge \text{WF}_t(A) \Rightarrow (P \rightsquigarrow Q)} \quad (\text{WF1})$$

In this rule, P and Q are state predicates, I is again an invariant, $[N]_t$ represents the next-state relation, and $\langle A \rangle_t$ is a “helpful action” [29] for which weak fairness is assumed. Again, all three premises of (WF1) are transition formulas. To see why the rule is correct, assume that σ is a behavior satisfying $\Box I \wedge \Box [N]_t \wedge \text{WF}_t(A)$, and that P holds of state σ_i . We have to show that Q holds of some σ_j with $j \geq i$. By the first premise, any successor of a state satisfying P has to satisfy P or Q , so P must hold for as long as Q has not been true. The third premise ensures that in all of these states, action $\langle A \rangle_t$ is enabled, and so the assumption of weak fairness ensures that eventually $\langle A \rangle_t$ occurs, unless Q has become true before, in which case we are done. Finally, by the second premise, any $\langle A \rangle_t$ -successor (which, by assumption, is in fact an $\langle N \wedge A \rangle_t$ -successor) of a state satisfying P must satisfy Q , which proves the claim.

For our running example, we can use rule (WF1) to prove that every message stored in the queue will eventually move closer to the head of the queue or even to the output channel. Formally, let the state predicate $at(k, x)$ be defined by

$$at(k, x) \triangleq k \in 1..Len(q) \wedge q[k] = x$$

We will use (WF1) to prove

$$FifoI \Rightarrow (at(k, x) \rightsquigarrow (out = x \vee at(k-1, x))) \quad (7)$$

where k and x are rigid variables. The following proof outline illustrates the application of rule (WF1), the lower-level steps relying on data axioms are omitted.

1. $at(k, x) \wedge [Next]_{vars} \Rightarrow at(k, x)' \vee out' = x \vee at(k-1, x)'$
 - 1.1. $at(k, x) \wedge m \in Message \wedge Enq(m) \Rightarrow at(k, x)'$
 - 1.2. $at(k, x) \wedge Deq \wedge k = 1 \Rightarrow out' = x$
 - 1.3. $at(k, x) \wedge Deq \wedge k > 1 \Rightarrow at(k-1, x)'$
 - 1.4. $at(k, x) \wedge vars' = vars \Rightarrow at(k, x)'$
 - 1.5. Q.E.D.

From steps 1.1–1.4 by the definitions of $Next$ and $at(k, x)$.

2. $at(k, x) \wedge \langle Deq \wedge Next \rangle_{vars} \Rightarrow out' = x \vee at(k-1, x)'$

Follows from steps 1.2 and 1.3 above.

3. $at(k, x) \Rightarrow \text{ENABLED } \langle Deq \rangle_{vars}$

For any k , $at(k, x)$ implies that $q \neq \langle \rangle$ and thus the enabledness condition.

However, rule (WF1) cannot be used to prove the stronger property that every input to the queue will eventually be dequeued,

$$\text{FifoI} \Rightarrow \forall m \in \text{Message} : \text{in} = m \rightsquigarrow \text{out} = m \quad (8)$$

because there is no single “helpful action”: the number of *Deq* actions necessary to produce the input element on the output channel depends on the length of the queue. Intuitively, the argument used to establish property (7) must be iterated. The following rule formalizes this idea as an induction over a well-founded relation (D, \succ) : a binary relation such that there does not exist an infinite descending chain $d_1 \succ d_2 \succ \dots$ of elements $d_i \in D$.

$$\frac{(D, \succ) \text{ is well-founded} \quad F \Rightarrow \forall d \in D : (G \rightsquigarrow (H \vee \exists e \in D : d \succ e \wedge G[e/d]))}{F \Rightarrow \forall d \in D : (G \rightsquigarrow H)} \quad (\text{LATTICE})$$

In this rule, d and e are rigid variables such that d does not occur in H and e does not occur in G . For convenience, we have stated rule (LATTICE) in a language of set theory where, in particular, $\forall x \in S : F$ abbreviates the formula $\forall x : x \in S \Rightarrow F$.

Unlike the premises of the rules considered so far, the second hypothesis of rule (LATTICE) is itself a temporal formula that requires that every occurrence of G , for any value $d \in D$, be followed either by an occurrence of H , or again by some G , for some smaller value e . Because the first hypothesis ensures that there cannot be an infinite descending chain of values in D , eventually H must become true. In principle, the hypothesis of well-foundedness can itself be expressed in TLA by asserting the validity of the formula

$$\begin{aligned} & \wedge \forall d \in D : \neg(d \succ d) \\ & \wedge \forall v : \Box(v \in D) \wedge \Box[v \succ v']_v \Rightarrow \Diamond\Box[\text{FALSE}]_v \end{aligned}$$

whose first conjunct expresses the irreflexivity of \succ and whose second conjunct asserts that any sequence of values in D that can only change by decreasing with respect to \succ must eventually become stationary. In fact, if this formula is valid over a given interpretation then \succ is interpreted by a well-founded relation. In system verification, well-foundedness is however usually considered as a “data axiom”.

Choosing $(\text{Nat}, >)$, the set of natural numbers with the standard “greater-than” relation as the well-founded domain, the proof of property (8) follows from property (7) and the invariant *Inv* defined in section 3.2 using rule (LATTICE).

Lamport [25] lists further (derived) rules for liveness properties, including introduction rules for proving formulas $\text{WF}_t(A)$ and $\text{SF}_t(A)$.

3.4 Simple temporal logic

The proof rules considered so far support the derivation of typical correctness properties of systems. In addition, TLA satisfies standard axioms and rules of linear-time

(STL1) $\frac{F}{\Box F}$	(STL4) $\Box(F \Rightarrow G) \Rightarrow (\Box F \Rightarrow \Box G)$
(STL2) $\Box F \Rightarrow F$	(STL5) $\Box(F \wedge G) \equiv (\Box F \wedge \Box G)$
(STL3) $\Box \Box F \equiv \Box F$	(STL6) $\Diamond \Box(F \wedge G) \equiv (\Diamond \Box F \wedge \Diamond \Box G)$

Fig. 2. Simple temporal logic.

temporal logic that are useful when preparing the application of verification rules. Figure 2 contains the axioms and rules of “simple temporal logic”, adapted from Lamport [25]. It can be shown that this is just a non-standard presentation of the modal logic S4.2 [19], implying that these laws by themselves characterize a modal accessibility relation for \Box that is reflexive, transitive, and locally convex (or confluent). The last condition asserts that for any state s and states t, u that are both accessible from s there is a state v that is accessible from t and u .

3.5 Quantifier rules

Although we have seen in section 2.4 that the semantics of quantification over flexible variables is non-standard, the elementary proof rules for quantifiers are those familiar from first-order logic:

$$\begin{array}{ll}
 F[c/x] \Rightarrow \exists x : F & (\exists I) \\
 \frac{F \Rightarrow G}{(\exists x : F) \Rightarrow G} & (\exists E) \\
 F[t/v] \Rightarrow \exists v : F & (\exists I) \\
 \frac{F \Rightarrow G}{(\exists v : F) \Rightarrow G} & (\exists E)
 \end{array}$$

In these rules, x is a rigid and v is a flexible variable. The elimination rules $(\exists E)$ and $(\exists E)$ require the usual proviso that the bound variable should not be free in formula G . In the introduction rules, t is a state function, while c is a constant function. Observe that if we allowed an arbitrary state function in rule $(\exists I)$, we could prove

$$\exists x : \Box(x = v) \quad (9)$$

for any state variable v from the premise $\Box(v = v)$, provable by (STL1). However, formula (9) asserts that v remains constant throughout a behavior, which can obviously not be valid.

Since existential quantification over flexible variables corresponds to hiding of state components, the rules $(\exists I)$ and $(\exists E)$ play a fundamental role in proofs of refinement for reactive systems. In this context, the “witness” t is often called a *refinement mapping* [2]. For example, the concatenation of the two low-level

queues provides a suitable refinement mapping to prove the validity of formula (1), which claimed that two FIFO queues in a row implement a FIFO queue, assuming interleaving of changes to the input and output channels.

Although the quantifier rules are standard, one should recall from section 2.2 that care has to be taken when substitutions are applied to formulas that are defined in terms of quantification. In particular, the formulas $WF_t(A)$ and $SF_t(A)$ contain the subformula $ENABLED \langle A \rangle_t$, and therefore, e.g., $WF_t(A)[e/v]$ need not be equivalent to the formula $WF_{t[e/v]}(A[e/v, e'/v'])$. For a more thorough discussion of this possible pitfall in system verification, see Lamport's original TLA paper [25].

Unfortunately, refinement mappings need not always exist. For example, $(\exists I)$ cannot be used to prove the valid TLA formula (excluding one-element universes)

$$\exists v : \Box \Diamond \langle TRUE \rangle_v \quad (10)$$

that asserts the existence of a flexible variable whose value changes infinitely often. (Such a variable could be used as an “oscillator”, triggering system transitions.) In fact, an attempt to prove (10) by rule $(\exists I)$ would require to exhibit a state function t whose value is certain to change infinitely often in any behavior. However, it is easy to show by induction on the syntax of state functions that for any t there exists a behavior such that the value of t remains constant forever.

An approach to solving this problem, advocated in [2], consists in adding *auxiliary variables* such as history and prophecy variables. Formally, this approach consists in adding special introduction rules for auxiliary variables. The proof of $G \Rightarrow \exists v : F$ is then reduced to first proving a formula of the form $G \Rightarrow \exists a : G_{aux}$ using a rule for auxiliary variables, and then use the rules $(\exists E)$ and $(\exists I)$ above to prove $G \wedge G_{aux} \Rightarrow \exists v : F$.

4 FORMALIZED MATHEMATICS: THE ADDED VALUE OF TLA^+

The definitions of the syntax and semantics of TLA in section 2 were generic in terms of an underlying language of predicate logic and its interpretation. TLA^+ instantiates this generic definition of TLA with a specific first-order language, namely Zermelo-Fr ankel set theory with choice. This adoption of a standard interpretation enables precise and unambiguous specifications of the “data structures” on which specifications are based; we have seen in the example proofs in section 3 that reasoning about the data accounts for most of the steps that need to be proved during system verification. TLA^+ also provides facilities for structuring a specification as a hierarchy of modules, for declaring parameters, and most importantly, for defining operators. These facilities are essential for writing actual specifications and must therefore be mastered by any user of TLA^+ . However, from the foundational point of view adopted in this paper, they are just syntactic sugar. We will therefore concentrate on the set-theoretic foundations, referring the reader to Lamport's book [27] for a detailed presentation of the language of TLA^+ .

4.1 Elementary data structures: basic set theory

The signature of the predicate logic underlying TLA^+ contains a single binary predicate symbol \in and no function symbols.⁴ Terms and formulas at the transition level are defined as indicated in section 2.2, with an extra term formation rule that defines $\text{CHOOSE } x : A$ to be a transition function whenever $x \in \mathcal{V}_E$ is a variable and A is an action.⁵ The occurrences of x in the term $\text{CHOOSE } x : A$ are bound. To this first-order language corresponds a set-theoretic interpretation: every TLA^+ value is a set. Moreover, \in is interpreted as set membership and the interpretation is equipped with an (unspecified) choice function ε mapping every non-empty collection C of values to some element $\varepsilon(C)$ of C , and mapping the empty collection to an arbitrary value. The interpretation of a term $\text{CHOOSE } x : P$ is defined as

$$\llbracket \text{CHOOSE } x : P \rrbracket_{s,t}^{\varepsilon} = \varepsilon(\{d \mid \llbracket P \rrbracket_{\alpha_{s,t}, \varepsilon(d/x)} = \mathbf{t}\})$$

This definition employs the choice function to return some value satisfying P provided there is some such value in the universe of set theory. We should remark that in this semantic clause, the choice function is applied to a collection that need not be a set (i.e., an element of the universe); in set-theoretic terminology, ε applies to classes and not just to sets. Because ε is a function, it produces the same value when applied to equal arguments. It follows that choice satisfies the laws

$$(\exists x : P) \equiv P[(\text{CHOOSE } x : P)/x] \quad (11)$$

$$(\forall x : (P \equiv Q)) \Rightarrow (\text{CHOOSE } x : P) = (\text{CHOOSE } x : Q) \quad (12)$$

TLA^+ avoids undefinedness by underspecification [17], so $\text{CHOOSE } x : P$ denotes a value even if no value satisfies P . To ensure that a term involving choice actually denotes the expected value, the existence of some set satisfying the characteristic predicate should be proven. If there is more than one such value, the expression is underspecified, and the user should be prepared to accept any of them. In particular, any properties will have to be established for all possible values. However, observe that for a given interpretation, choice is deterministic and not “monotone”: no relationship can be established between $\text{CHOOSE } x : P$ and $\text{CHOOSE } x : Q$ even when $P \Rightarrow Q$ is valid (unless P and Q are actually equivalent). Therefore, whenever some specification Spec contains an underspecified application of choice, any refinement Ref is bound to make the same choices in order to prove $\text{Ref} \Rightarrow \text{Spec}$; this situation is quite different from non-determinism where implementations may narrow the set of allowed values.

In the following, we will freely use many abbreviations defined by TLA^+ . For example, $\exists x, y \in S : P$ abbreviates $\exists x : \exists y : x \in S \wedge y \in S \wedge P$, and similar

⁴ Once again, our presentation deviates somewhat from Lamport [27] who prefers to treat all subsequent constructs on an equal footing rather than distinguishing between basic and derived operators.

⁵ Temporal formulas are defined as indicated in section 2.3; in particular, CHOOSE is never applied to a temporal formula.

notation applies to \forall and CHOOSE. Local declarations are written as LET _ IN _, and IF _ THEN _ ELSE _ is used for conditional expressions.

union	$UNION\ S \triangleq CHOOSE\ M : \forall x : (x \in M \equiv \exists T \in S : x \in T)$
binary union	$S \cup T \triangleq UNION\ \{S, T\}$
subset	$S \subseteq T \triangleq \forall x : (x \in S \Rightarrow x \in T)$
powerset	$SUBSET\ S \triangleq CHOOSE\ M : \forall x : (x \in M \equiv x \subseteq S)$
comprehension 1	$\{x \in S : P\} \triangleq CHOOSE\ M : \forall x : (x \in M \equiv x \in S \wedge P)$
comprehension 2	$\{t : x \in S\} \triangleq CHOOSE\ M : \forall y : (y \in M \equiv \exists x \in S : y = t)$

Table 1. Basic set-theoretic operators.

From membership and choice, one can build up the conventional language of mathematics [28], and this is the foundation for the expressiveness of TLA^+ . Table 1 lists some of the basic set-theoretic constructs of TLA^+ ; we write

$$\{e_1, \dots, e_n\} \triangleq CHOOSE\ S : \forall x : (x \in S \equiv x = e_1 \vee \dots \vee x = e_n)$$

to denote set enumeration and assume the additional bound variables in the defining expressions of table 1 to be chosen such that no variable clashes occur. The two comprehension schemes act as binders for variable x , which must not have free occurrences in S . The existence of the sets defined in terms of choice can be justified from the axioms of Zermelo-Fr nkel set theory [37], which provide the deductive counterpart to the semantics underlying TLA^+ . However, it is well-known that without proper care, set theory is prone to paradoxes. For example, the expression

$$CHOOSE\ S : \forall x : (x \in S \equiv x \notin x)$$

is a well-formed constant formula of TLA^+ , but the existence of a set S containing precisely those sets that do not contain themselves would lead to the contradiction that $S \in S$ iff $S \notin S$; this is of course Russell's paradox. Intuitively, S is "too big" to be a set. More precisely, the universe of set theory does not contain collections that are in bijection with the collection of all sets. Therefore, when evaluating the above TLA^+ expression, the choice function is applied to the empty collection, and the result depends on the underlying interpretation. Perhaps unexpectedly, we can however infer from (12) that

$$(CHOOSE\ S : \forall x : (x \in S \equiv x \notin x)) = (CHOOSE\ x : x \in \{\})$$

Similarly, a generalized intersection operator dual to the union operator of table 1 cannot be sensibly defined, because the intersection of the empty set would have to produce the set of all sets, which we know cannot exist.

On the positive side, we have exploited the fact that no set can contain all values in the definition

$$NoMsg \triangleq CHOOSE\ x : x \notin Message$$

that appears in figure 1(b) because *NoMsg* is guaranteed to denote some value that is not contained in *Message*. If a later refinement wanted to fix a specific “null” message value $\text{null} \notin \text{Message}$, it could do so by restricting the class of admissible interpretations via an assumption of the form

$$\text{ASSUME } (\text{CHOOSE } x : x \notin \text{Message}) = \text{null}$$

Because all properties established of the original specification hold for all possible choices of *NoMsg*, they will continue to hold for this restricted choice.

4.2 More data structures

Some sets can conveniently be interpreted as functions. A traditional approach, followed in Z and B [8, 36], is to construct functions via products and relations. TLA⁺ does not prescribe any concrete construction of functions. The set of functions whose domain equals *S* and whose codomain is a subset of *T* is denoted by $[S \rightarrow T]$, the domain of a function *f* is denoted by $\text{DOMAIN } f$, and the application of function *f* to an expression *e* is written as $f[e]$. The expression $[x \in S \mapsto e]$ denotes the function with domain *S* that maps any $x \in S$ to *e*; again, the variable *x* must not occur in *S* and is bound by the function constructor. Thus, any function *f* obeys the law

$$f = [x \in \text{DOMAIN } f \mapsto f[x]] \quad (13)$$

and this equation can in fact serve as a characteristic predicate for functional values. TLA⁺ introduces a notation for overriding a function at a certain argument position (a similar concept underlies Gurevich’s ASM notation [12]). Formally,

$$[f \text{ EXCEPT } ![t] = u] \triangleq [x \in \text{DOMAIN } f \mapsto \text{IF } x = t \text{ THEN } u \text{ ELSE } f[x]]$$

where *x* is a fresh variable.

Combining choice, sets, and function notation, one obtains an expressive language for defining mathematical structures. For example, the standard TLA⁺ module introducing natural numbers defines them as an arbitrary set with constant zero and successor function satisfying the usual Peano axioms [27, p. 345], and Lamport goes on to similarly define the integers and the real numbers, ensuring that the basic arithmetic operations agree rather than having to overload the operation symbols.

Recursive definitions can be introduced in terms of choice, e.g.

$$\text{factorial} \triangleq \text{CHOOSE } f : f = [n \in \text{Nat} \mapsto \text{IF } n = 0 \text{ THEN } 1 \text{ ELSE } n * f[n - 1]]$$

which TLA⁺, using some syntactic sugar, allows to write even more concisely as

$$\text{factorial}[n \in \text{Nat}] \triangleq \text{IF } n = 0 \text{ THEN } 1 \text{ ELSE } n * \text{factorial}[n - 1]$$

Of course, as with any construction based on choice, such a definition should be justified by proving the existence of a function that satisfies the recursive equation. Unlike standard semantics of programming languages, TLA⁺ does not commit to the least fixed point of a recursively defined function in cases where there are several solutions.

Tuples are represented in TLA⁺ as functions:

$$\langle t_1, \dots, t_n \rangle \triangleq [i \in 1..n \mapsto \text{IF } i = 1 \text{ THEN } t_1 \dots \text{ ELSE } t_n]$$

where $1..n$ denotes the set $\{j \in \text{Nat} : 1 \leq j \wedge j \leq n\}$ (and i is a “fresh” variable), and selection of the i -th element is just function application. Strings are defined as tuples of characters, and records are represented as functions whose domains are finite sets of strings. The update operation on functions can thus be applied to tuples and records as well. For record selection and update, the concrete syntax allows to write, for example, *acct.balance* instead of *acct*["balance"].

<i>Seq</i> (<i>S</i>)	\triangleq	UNION $\{[1..n] \rightarrow S : n \in \text{Nat}\}$
<i>Len</i> (<i>s</i>)	\triangleq	CHOOSE $n \in \text{Nat} : \text{DOMAIN } s = 1..n$
<i>Head</i> (<i>s</i>)	\triangleq	<i>s</i> [1]
<i>Tail</i> (<i>s</i>)	\triangleq	$[i \in 1..Len(s) - 1 \mapsto s[i + 1]]$
<i>s</i> \circ <i>t</i>	\triangleq	$[i \in 1..Len(s) + Len(t) \mapsto$ IF $i \leq Len(s)$ THEN $s[i]$ ELSE $t[i - Len(s)]]$
<i>Append</i> (<i>s</i> , <i>e</i>)	\triangleq	<i>s</i> \circ $\langle e \rangle$

Fig. 3. Finite sequences.

The standard TLA⁺ module *Sequences* imported by the specification of the FIFO queue in figure 1(b) represents finite sequences as tuples. The definitions of the standard operations, some of which are shown in figure 3, is therefore quite simple. However, this simplicity can sometimes be deceptive. For example, these definitions do not reveal that the *Head* and *Tail* operations are “partial”. They should be validated by proving the expected properties, such as

$$\forall s \in Seq(S) : Len(s) \geq 1 \Rightarrow s = \langle Head(s) \rangle \circ Tail(s)$$

5 CONCLUSIONS

The design of software systems requires a combination of ingenuity and careful engineering. While there is no substitute for intuition, the correctness of a proposed solution can be checked by precise reasoning over a suitable model, and this is the realm of logics and (formalized) mathematics. The rôle of a formalism is to *help* the user in the difficult and important activity of writing and analysing formal models. TLA⁺ builds on the experience of classical mathematics and adds just a thin layer of

temporal logic in order to describe executions as sets of traces. A distinctive feature of TLA is its attention to refinement and composition, reflected in the concept of stuttering invariance.

Whereas the expressiveness of TLA^+ undoubtedly helps in writing concise, high-level models of systems, one may wonder whether it lends itself as well to the analysis of these models. For example, we have pointed out several times the need to prove conditions of “well-definedness” related to the use of the choice operator. On the other hand, Zermelo-Fränkel set theory with choice is probably the most widely used foundation of classical mathematics, and there are well-known idioms, such as primitive-recursive definitions, that ensure well-definedness. For the specification of reactive systems, TLA adds some proper idioms that control the delicate interplay between temporal operators, e.g. in order to ensure that a specification is machine closed [3].

Deductive verification of TLA^+ specifications can be supported by proof assistants, and in fact several encodings of TLA in the logical frameworks of different theorem provers have been proposed [15, 20, 31], although no prover has yet been designed to support full TLA^+ . Perhaps more surprisingly, there has been much recent activity on developing a toolset based on the TLC model checker and simulator to aid in validating and debugging TLA^+ models [39], and this toolset has been applied in industrial development projects. Obviously, model checking is possible only for a sublanguage of TLA^+ , but interestingly, most real-world specifications are either written in this sublanguage or can be translated into it using minor transformations. The modeling language of TLC is still much more expressive than that of most other model checkers and therefore helps users to write concise system specifications.

Acknowledgements. I am grateful to Leslie Lamport for providing the subject of this article, for his encouragement of this work, and for detailed comments on a draft version. Parts of this paper have been adapted from an earlier paper on TLA, written with Martín Abadi [6]. Helpful comments by Júlia Zappe and by the anonymous referees are gratefully acknowledged.

REFERENCES

- [1] Martín Abadi. An axiomatization of Lamport’s Temporal Logic of Actions. In Jos C. M. Baeten and Jan W. Klop, editors, *CONCUR ’90, Theories of Concurrency: Unification and Extension*, volume 458 of *Lecture Notes in Computer Science*, pages 57–69, Berlin, 1990. Springer-Verlag.
- [2] MARTÍN ABADI AND LESLIE LAMPORT. THE EXISTENCE OF REFINEMENT MAPPINGS. *Theoretical Computer Science*, 81(2):253–284, May 1991.
- [3] MARTÍN ABADI AND LESLIE LAMPORT. AN OLD-FASHIONED RECIPE FOR REAL TIME. *ACM Transactions on Programming Languages and Systems*, 16(5):1543–1571, September 1994.

- [4] MARTÍN ABADI AND LESLIE LAMPORT. CONJOINING SPECIFICATIONS. *ACM Transactions on Programming Languages and Systems*, 17(3):507–534, May 1995.
- [5] Martín Abadi and Stephan Merz. An abstract account of composition. In Jiří Wiedermann and Petr Hajek, editors, *Mathematical Foundations of Computer Science*, volume 969 of *Lecture Notes in Computer Science*, pages 499–508, Prague, Czech Republic, 1995. Springer-Verlag.
- [6] Martín Abadi and Stephan Merz. On TLA as a logic. In Manfred Broy, editor, *Deductive Program Design*, NATO ASI series F, pages 235–272. Springer-Verlag, 1996.
- [7] MARTÍN ABADI AND GORDON PLOTKIN. A LOGICAL VIEW OF COMPOSITION. *Theoretical Computer Science*, 114(1):3–30, June 1993.
- [8] J.-R. Abrial. *The B-Book: Assigning Programs to Meanings*. Cambridge University Press, 1996.
- [9] J.-R. Abrial. Extending B without changing it (for developing distributed systems). In H. Habrias, editor, *1st Conference on the B method*, pages 169–190. IRIN Institut de recherche en informatique de Nantes, 1996.
- [10] BOWEN ALPERN AND FRED B. SCHNEIDER. DEFINING LIVENESS. *Information Processing Letters*, 21(4):181–185, October 1985.
- [11] R. Back and J. von Wright. *Refinement calculus—A systematic introduction*. Springer-Verlag, 1998.
- [12] Egon Börger and Robert Stärk. *Abstract State Machines: A Method for High-Level System Design and Analysis*. Springer-Verlag, 2003.
- [13] Edmund M. Clarke, Orna Grumberg, and Doron Peled. *Model Checking*. MIT Press, Cambridge, Mass., 1999.
- [14] W.-P. de Roever, H. Langmaack, and A. Pnueli, editors. *Compositionality: The Significant Difference*, volume 1536 of *Lecture Notes in Computer Science*. Springer-Verlag, 1998.
- [15] Urban Engberg, Peter Gronning, and Leslie Lamport. Mechanical verification of concurrent systems with TLA. In *Fourth Intl. Conf. Computer-Aided Verification (CAV '92)*, volume 663 of *Lecture Notes in Computer Science*, pages 44–55. Springer-Verlag, 1992.
- [16] M. C. J. Gordon. Mechanizing programming logics in higher-order logic. In G. Birtwistle and P. A. Subrahmanyam, editors, *Current Trends in Hardware Verification and Automated Theorem Proving*. Springer-Verlag, 1989.
- [17] David Gries and Fred B. Schneider. Avoiding the undefined by underspecification. In Jan van Leeuwen, editor, *Computer Science Today: Recent Trends and Developments*, volume 1000 of *Lecture Notes in Computer Science*, pages 366–373. Springer-Verlag, New York, N.Y., 1995.
- [18] C. A. R. Hoare. A theory of conjunction and concurrency. In *Parallel Architectures (Parbase '90)*, pages 18–27. IEEE Computer Security Press, 1991.
- [19] G. E. Hughes and M. J. Cresswell. *A New Introduction to Modal Logic*. Routledge, London, 1968.
- [20] Sara Kalvala. A formulation of TLA in Isabelle. Available at <ftp://ftp.dcs.warwick.ac.uk/people/Sara.Kalvala/tla.dvi>, March 1995.

- [21] H. W. Kamp. *Tense Logic and the Theory of Linear Order*. PhD thesis, Univ. of California at Los Angeles, 1968.
- [22] Leslie Lamport. The TLA home page. <http://www.research.microsoft.com/users/lamport/tla/tla.html>.
- [23] LESLIE LAMPORT. PROVING THE CORRECTNESS OF MULTIPROCESS PROGRAMS. *IEEE Transactions on Software Engineering*, SE-3(2):125–143, March 1977.
- [24] Leslie Lamport. What good is temporal logic? In R. E. A. Mason, editor, *Information Processing 83: Proceedings of the IFIP 9th World Congress*, pages 657–668, Paris, September 1983. IFIP, North-Holland.
- [25] LESLIE LAMPORT. THE TEMPORAL LOGIC OF ACTIONS. *ACM Transactions on Programming Languages and Systems*, 16(3):872–923, May 1994.
- [26] LESLIE LAMPORT. HOW TO WRITE A PROOF. *American Mathematical Monthly*, 102(7):600–608, 1995.
- [27] Leslie Lamport. *Specifying Systems*. Addison-Wesley, Boston, Mass., 2002.
- [28] A. C. Leisenring. *Mathematical Logic and Hilbert's ϵ -Symbol*. University Mathematical Series. Macdonald & Co. Ltd., London, U.K., 1969.
- [29] Zohar Manna and Amir Pnueli. Verification of concurrent programs: the temporal framework. In R.S. Boyer and J.S. Moore, editors, *The Correctness Problem in Computer Science*, pages 215–273. Academic Press, London, 1982.
- [30] Zohar Manna and Amir Pnueli. *The temporal logic of reactive and concurrent systems—Specification*. Springer-Verlag, New York, 1992.
- [31] Stephan Merz. Isabelle/TLA. Available on the WWW at <http://isabelle.in.tum.de/library/HOL/TLA>, 1997. Revised 1999.
- [32] Stephan Merz. A more complete TLA. In J.M. Wing, J. Woodcock, and J. Davies, editors, *FM'99: World Congress on Formal Methods*, volume 1709 of *Lecture Notes in Computer Science*, pages 1226–1244, Toulouse, France, 1999. Springer-Verlag.
- [33] Amir Pnueli. The temporal logic of programs. In *Proceedings of the 18th Annual Symposium on the Foundations of Computer Science*, pages 46–57. IEEE, 1977.
- [34] Arthur N. Prior. *Past, Present and Future*. Clarendon Press, Oxford, U.K., 1967.
- [35] A. P. SISTLA, E. M. CLARKE, N. FRANCEZ, AND Y. GUREVICH. CAN MESSAGE BUFFERS BE CHARACTERIZED IN LINEAR TEMPORAL LOGIC? *Information and Control*, 63:88–112, 1984.
- [36] M. Spivey. *The Z Notation: A Reference Manual*. Prentice Hall, 1992.
- [37] Patrick Suppes. *Axiomatic Set Theory*. Dover Publications, 1972.
- [38] Moshe Vardi. Branching vs. linear time—final showdown. In T. Margaria and W. Yi, editors, *Tools and Algorithms for the Construction and Analysis of Systems (TACAS 2001)*, volume 2031 of *Lecture Notes in Computer Science*, pages 1–22, Genova, Italy, 2001. Springer-Verlag. See <http://www.cs.rice.edu/~vardi/papers/> for more recent versions of this paper.
- [39] Yuan Yu, Panagiotis Manolios, and Leslie Lamport. Model checking TLA+ specifications. In L. Pierre and T. Kropf, editors, *Correct Hardware Design and Verification Methods (CHARME'99)*, volume 1703 of *Lecture Notes in Computer Science*, pages 54–66, Bad Herrenalb, Germany, 1999. Springer-Verlag.

EXHIBIT B

Emacs

From Wikipedia, the free encyclopedia

This article is about the text editor. For the unrelated Apple Macintosh computer model, see eMac.

Emacs is a class of text editors, possessing an extensive set of features, that are popular with computer programmers and other technically proficient computer users.

GNU Emacs, a part of the GNU project, is under active development and is the most popular version. The GNU Emacs manual describes it as "the extensible, customizable, self-documenting, real-time display editor." It is also the most portable and ported of the implementations of Emacs. As of 2007, the latest stable release of GNU Emacs is version 21.4.

The original EMACS, a set of *Editor MACroS* for the TECO editor, was written in 1975 by Richard Stallman, initially put together with Guy Steele. It was inspired by the ideas of TECMAC and TMACS, a pair of TECO-macro editors written by Guy Steele, Dave Moon, Richard Greenblatt, Charles Frankston, and others. Many versions of Emacs have appeared over the years, but nowadays there are two that are commonly used: GNU Emacs, started by Richard Stallman in 1984 and still maintained by him, and XEmacs, a fork of GNU Emacs which was started in 1991 and has remained mostly compatible. Both use a powerful extension language, Emacs Lisp, that allows them to handle tasks ranging from writing and compiling computer programs to browsing the web.

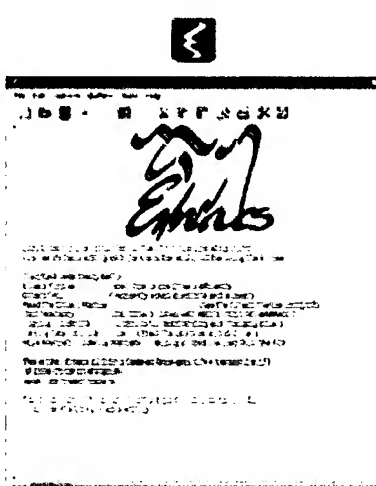
Some people make a distinction between the capitalized word "Emacs", used to refer to editors derived from versions created by Richard Stallman, and the lower-case word "emacs", which is used to refer to the large number of independent emacs reimplementations. The word "emacs" is often pluralized as **emacsen** by analogy with "oxen". For example, Debian's compatible Emacs package is named `emacs-en-common`. The only plural given by the *Collins English Dictionary* is *emacsen*.^[1]

In Unix culture, Emacs is one of the two main contenders in the traditional editor wars, the other being vi.

Contents

- 1 History
 - 1.1 Other emacsen
 - 1.2 GNU Emacs
 - 1.3 XEmacs
 - 1.4 Other implementations

Emacs



The GNU Emacs 22 interface, running in a graphical environment

Maintainer:	GNU Project
Stable release:	21.4 (Feb 6, 2005) [+/-]
Preview release:	22 (December 4, 2005) [+/-]
OS:	Cross-platform
Available language (s):	English only
Use:	Text editor
License:	GPL
Website:	www.gnu.org/software/emacs/

- 2 Licensing
- 3 Features
 - 3.1 Platforms
 - 3.2 Editing modes
 - 3.3 Customization
 - 3.4 Documentation
 - 3.5 Internationalization
- 4 License
- 5 Using Emacs
 - 5.1 Commands
 - 5.2 Minibuffer
 - 5.3 File management and display
 - 5.4 Emacs Pinky
 - 5.5 Distractions
- 6 See also
- 7 References
- 8 Bibliography
- 9 External links

History

Emacs began life at the MIT AI Lab during the 1970s. Before its introduction, the default editor on the Incompatible Timesharing System (ITS), the operating system on the AI Lab's PDP-6 and PDP-10 computers, was a line editor known as TECO. Unlike modern text editors, TECO treated typing, editing, and document display as separate modes, as the later vi would. Typing characters into TECO did not place those characters directly into a document; one had to write a series of instructions in the TECO command language telling it to enter the required characters, during which time the edited text was not displayed on the screen. This behavior is similar to the program ed, which is still in use.

Richard Stallman visited the Stanford AI Lab in 1972 or 1974 and saw the lab's "E" editor, written by Fred Wright. The editor had an intuitive WYSIWYG behavior as is used almost universally by modern text editors. Impressed by this feature, Stallman returned to MIT where Carl Mikkelsen, one of the hackers at the AI Lab, had added a display-editing mode called "Control-R" to TECO, allowing the screen display to be updated each time the user entered a keystroke. Stallman reimplemented this mode to run efficiently, then added a macro feature to the TECO display-editing mode, allowing the user to redefine any keystroke to run a TECO program.

Another feature of E which was lacking in TECO was random-access editing. Since TECO's original implementation was designed for editing paper tape on the PDP-1, it was a page-sequential editor. Typical editing could only be performed on one page at a time, in the order that the pages appeared in the file. To provide random access in Emacs, Stallman elected not to adopt E's approach of structuring the file for page-random access on disk, but instead modified TECO to handle large buffers more efficiently, and then changed its file management philosophy to read, edit, and write the entire file as a single buffer. Almost all modern editors use this approach.

The new version of TECO was instantly popular at the AI Lab, and soon there accumulated a large collection of custom macros, whose names often ended in "MAC" or "MACS", which stood for "macros". Two years later, Guy Steele took on the project of unifying the overly diverse keyboard command sets into a single set. After one night of joint hacking by Steele and Stallman, the latter finished the implementation, which included facilities for extending and documenting the new macro set. The resulting system was called EMACS, which stood for "Editing MACroS". An alternate version is that EMACS stood for "E with MACroS", a dig at E's lack of a macro capability. According to Stallman, he picked the name Emacs "because <E> was not in use as an abbreviation on ITS at the time." It has also

been pointed out that "Emack & Bolio's" was the name of a popular ice cream store in Boston, within walking distance of MIT. A text-formatting program used on ITS was later named BOLIO by Dave Moon, who frequented that store. However, Stallman did not like that ice cream, and did not even know of it when choosing the name "Emacs"; this ignorance is the basis of a Hacker koan, *Emacs and Bolio*).

Stallman realized the danger of too much customization and de-facto forking and set certain conditions for usage. He later wrote:

"EMACS was distributed on a basis of communal sharing, which means all improvements must be given back to me to be incorporated and distributed."

The original Emacs, like TECO, ran only on the PDP line. Its behavior was different enough from TECO to be considered a text editor in its own right. It quickly became the standard editing program on ITS. It was also ported from ITS to the Tenex and TOPS-20 operating systems by Michael McMahon, but not Unix, initially. Other contributors to early versions of Emacs include Kent Pitman, Earl Killian, and Eugene Ciccarelli.

Other emacsen

Many Emacs-like editors were written in the following years for other computer systems, including SINE (Sine is not EMACS), EINE ("EINE Is Not EMACS") and ZWEI ("ZWEI Was EINE Initially", for the Lisp machine), which were written by Michael McMahon and Daniel Weinreb ("eine" and "zwei" mean "one" and "two" in German, respectively). In 1978, Bernard Greenberg wrote Multics Emacs at Honeywell's Cambridge Information Systems Lab, which was the first version to fully embrace Lisp as its extension language. Emacs (including GNU Emacs) later adopted using Lisp as the editor's extension language.

The first Emacs-like editor to run on Unix was Gosling Emacs, written in 1981 by James Gosling (who later invented NeWS and the Java programming language). It was written in C and, notably, used a language with Lisp-like syntax known as Mocklisp as an extension language. In 1984 it was proprietary software.

GNU Emacs

In 1984, Stallman began working on GNU Emacs to produce a free software alternative to Gosling Emacs; initially it was based on Gosling Emacs, but Stallman replaced the Mocklisp interpreter at its heart with a true Lisp interpreter, which entailed replacing nearly all of the code. It became the first program released by the nascent GNU project. GNU Emacs is written in C and provides Emacs Lisp (itself implemented in C) as an extension language. The first widely distributed version of GNU Emacs was 15.34, which appeared in 1985. (Versions 2 to 12 never existed. Earlier versions of GNU Emacs had been numbered "1.x.x", but sometime after version 1.12 the decision was made to drop the "1", as it was thought the major number would never change. Version 13, the first public release, was made on March 20, 1985.)

Like Gosling Emacs, GNU Emacs ran on Unix; however, GNU Emacs had more features, in particular a full-featured Lisp as extension language. As a result, it soon replaced Gosling Emacs as the *de facto* Emacs editor on Unix.

Until 1999, GNU Emacs development was relatively closed, to the point where it was used as an example of the "Cathedral" development style in *The Cathedral and the Bazaar*. The project has since adopted a public development mailing list and anonymous CVS access. Development takes place in a single CVS trunk, which is at version 22.0.96. The current maintainer is Richard Stallman.

XEmacs

Beginning in 1991, Lucid Emacs was developed by Jamie Zawinski and others at Lucid Inc., based on an early alpha version of GNU Emacs 19. The codebases soon diverged, and the separate development teams gave up^[2] trying to merge them back into a single program. This was one of the most famous early forks of a free software program. Lucid Emacs has since been renamed XEmacs; it and GNU Emacs remain the two most popular varieties in use today.

Other implementations

GNU Emacs was initially targeted at computers with a 32-bit flat address space, and at least 1 MiB of RAM, at a time where such computers were considered high end. This left an opening for smaller reimplementations. Some noteworthy ones are listed here:

- MicroEMACS, a very portable implementation originally written by Dave Conroy and further developed by Daniel Lawrence, which exists in many variations. The editor used by Linus Torvalds.^[3]
- MG, originally called MicroGNUEmacs, an offshoot of MicroEMACS intended to more closely resemble GNU Emacs. Now installed by default on OpenBSD.
- JOVE (Jonathan's Own Version of Emacs), a non-programmable Emacs implementation for UNIX-like systems by Jonathan Payne.
- Freemacs, a DOS version with an extension language based on text macro expansion, all within the original 64 KiB flat memory limit.
- Meadow ^[4] is an Emacs variant originating from Japan that is designed to operate under Windows. The focus of Meadow is to provide multi-lingual support.
- MINCE (MINCE Is Not Complete Emacs), a version for CP/M from Mark of the Unicorn. MINCE evolved into Final Word, which eventually became the Sprint word processor from Borland.
- SXEmacs ^[1] is a fork of XEmacs 21.4.16 led by former XEmacs developer Steve Youngs. It aims to integrate Emacs into the X Windows systems, it includes MP3 player and productivity software.
- Zile

Licensing

For GNU Emacs (and GNU packages in general), it remains policy to accept significant code contributions only if the copyright holder executes a suitable disclaimer or assignment of their copyright interest, although one exception was made to this policy for the MULE (MULtilingual Extension, which handles Unicode and more advanced methods of dealing with other languages' scripts) code ^[2] since the copyright holder is the Japanese government and copyright assignment was not possible. This does not apply to extremely minor code contributions or bug fixes. There is no strict definition of minor, but as a guideline less than 10 lines of code is considered minor. This policy is intended to facilitate copyright enforcement, so that the FSF can defend the software in a court case if one arises.

Features

Emacs is a powerful and versatile text editor. It is primarily a *text* editor, not a word processor; it is geared toward manipulating pieces of text, rather than manipulating the font of the characters or printing documents (though Emacs can do these as well). Emacs provides commands to manipulate words and paragraphs (deleting them, moving them, moving through them, and so forth), syntax highlighting for making source code easier to read, and "keyboard macros" for performing arbitrary batches of editing commands defined by the user.

Almost all of the functionality in the editor, ranging from basic editing operations such as the insertion of characters into a document to the configuration of the user interface, is controlled by a dialect of the Lisp programming language

known as Emacs Lisp. This unique and unusual design provides many of the features found in Emacs. In this Lisp environment, variables and even entire functions can be modified on the fly, without having to recompile or even restart the editor. As a result, the behavior of Emacs can be modified almost without limit, either directly by the user, or by loading bodies of Emacs Lisp code known variously as "libraries", "packages", or "extensions".

Emacs contains a large number of Emacs Lisp libraries, and more "third-party" libraries can be found on the Internet. Many libraries implement computer programming aids, reflecting Emacs' popularity among programmers. Emacs can be used as an Integrated Development Environment (IDE), allowing programmers to edit, compile, and debug their code within a single interface. Other libraries perform more unusual functions. A few examples are listed below:

- Calc, a powerful RPN numerical calculator
- Calendar-mode, for keeping appointment calendars and diaries
- Doctor, an implementation of ELIZA that performs basic Rogerian psychotherapy
- Dunnet, a text adventure
- Ediff, for working with diff files interactively.
- Emerge, for comparing files and combining them
- Emacs/W3, a web browser
- ERC, an IRC client
- Gnus, a full-featured newsreader and email client
- MULE, MultiLingual extensions to Emacs, allowing editing text written in multiple languages, somewhat analogous to Unicode
- Info, an online help-browser
- *Emacs-wiki*, LISP-based wiki software for Emacs
- Planner, a personal information manager for Emacs
- Tetris
- Pong

The downside to Emacs' Lisp-based design is a performance overhead resulting from loading and interpreting the Lisp code. On the systems in which Emacs was first implemented, Emacs was often noticeably slower than rival text editors. Several joke backronyms allude to this: *Eight Megabytes And Constantly Swapping* (from the days when eight megabytes was a lot of memory), *Emacs Makes A Computer Slow, Eventually Mallocs All Computer Storage*, and *Eventually Makes All Computers Sick*. However, modern computers are fast enough that Emacs is seldom felt to be slow. In fact, Emacs starts up more quickly than most modern word processors. Other joke backronyms describe the user interface: *Escape Meta Alt Control Shift*.

Platforms

Emacs is one of the most ported non-trivial computer programs. It runs on a wide variety of operating systems, including most Unix-like systems (GNU/Linux, the various BSDs, Solaris, AIX, IRIX, Mac OS X,^{[5][6]} etc.), MS-DOS, Microsoft Windows^{[7][8][9]} and OpenVMS. Unix systems, both free and proprietary, frequently provide Emacs bundled with the operating system.

Emacs runs on both text terminals and graphical user interface (GUI) environments. On Unix-like operating systems, Emacs uses the X Window System to produce its GUI, either directly or using a "widget toolkit" such as Motif, LessTif, or GTK+. Emacs can also use the native graphical systems of Mac OS X (using the Carbon interface) and Microsoft Windows. The graphical interface provides menubars, toolbars, scrollbars, and context menus.

Editing modes

Emacs adapts its behavior to the type of text it is editing by entering editing modes called "major modes". Major modes are defined for ordinary text files, source code for many programming languages, HTML documents, TeX and LaTeX documents, and many other types of text. Each major mode tweaks certain Emacs Lisp variables to make Emacs behave more conveniently for the particular type of text. In particular, they usually implement syntax highlighting, using different fonts or colors to display keywords, comments, and so forth. Major modes also provide special editing commands; for example, major modes for programming languages usually define commands to jump to the beginning and the end of a function.

The behavior of Emacs can be further customized using "minor modes". While only one major mode can be associated with a buffer at a time, multiple minor modes can be simultaneously active. For example, the major mode for the C programming language defines a different minor mode for each of the popular indent styles.

Customization

Emacs can be customized to suit individual needs. There are three primary ways to customize Emacs. The first is the *customize* extension, which allows the user to set common customization variables, such as the colour scheme, using a graphical interface. This is intended for Emacs beginners who do not want to work with Emacs Lisp code.

The second is to collect keystrokes into macros and replay them to automate complex, repetitive tasks. This is often done on an ad-hoc basis and each macro discarded after use, although macros can be saved and invoked later.

The third method for customizing Emacs is using Emacs Lisp. Usually, user-supplied Emacs Lisp code is stored in a file called `.emacs`, which is loaded when Emacs starts up. The `.emacs` file is often used to set variables and key bindings different from the default setting, and to define new commands that the user finds convenient. Many advanced users have `.emacs` files hundreds of lines long, with idiosyncratic customizations that cause Emacs to diverge wildly from the default behavior.

If a body of Emacs Lisp code is generally useful, it is often packaged as a library and distributed to other users. Many such third-party libraries can be found on the Internet; for example, there is a library called `wikipedia-mode` for editing Wikipedia articles. There is even a Usenet newsgroup, `gnu.emacs.sources`, which is used for posting new libraries. Some third-party libraries eventually make their way into Emacs, thus becoming a "standard" library.

Documentation

The first Emacs included an innovative *help* library that can display the documentation for every single command, variable, and internal function. (It may have originated this technique.) Because of this, Emacs was described as "self-documenting". (This term does not mean that Emacs writes its own documentation, but rather that it presents its own documentation to the user.) This feature makes Emacs' documentation very accessible. For example, the user can find out about the command bound to a particular keystroke simply by entering `C-h k` (which runs the command `describe-key`), followed by the keystroke. Each function included a documentation string, specifically to be used for showing to the user on request. The practice of giving functions documentation strings subsequently spread to various programming languages such as Lisp and Java.

The Emacs *help* system is useful not only for beginners, but also for advanced users writing Emacs Lisp code. If the documentation for a function or variable is not enough, the *help* system can be used to browse the Emacs Lisp source code for both built-in libraries and installed third-party libraries. It is therefore very convenient to program in Emacs Lisp using Emacs itself.

Apart from the built-in documentation, Emacs has an unusually long, detailed and well-written manual. An electronic copy of the *GNU Emacs Manual*, written by Richard Stallman, is included with GNU Emacs and can be viewed with the built-in Info browser. XEmacs has a similar manual, which forked from the GNU Emacs Manual at the same time as the XEmacs software. Two other manuals, the *Emacs Lisp Reference Manual* by Bill Lewis, Richard Stallman, and Dan Laliberte, and *Programming in Emacs Lisp* by Robert Chassell, are also included. Apart from the electronic versions, all three manuals are also available in book form, published by the Free Software Foundation.

Emacs also has a built-in tutorial. When Emacs is started with no file to edit, it displays instructions for performing simple editing commands and invoking the tutorial.

Internationalization

Emacs supports the editing of text written in many human languages. There is support for many alphabets, scripts, writing systems, and cultural conventions. Emacs provides spell checking for many languages by calling external programs such as *ispell*. Many encoding systems, including UTF-8, are supported. XEmacs version 21.5 has partial Unicode support. Emacs 21.4 has similar support; Emacs 22 will be better. All of these efforts use an Emacs-specific encoding internally, necessitating conversion upon load and save. UTF-8 will become the Emacs-internal encoding in some later version of XEmacs 21.5, and likely in Emacs 23.

However, the Emacs user interface is in English, and has not been translated into any other language, with the exception of the beginners' tutorial.

For visually impaired and blind users, there is a subsystem called *Emacspeak* which allows the editor to be used through audio feedback only.

License

The source code, including both the C and Emacs Lisp components, is freely available for examination, modification, and redistribution, under the terms of the GNU General Public License (GPL). Older versions of the GNU Emacs documentation were released under an ad-hoc license which required the inclusion of certain text in any modified copy. In the GNU Emacs user's manual, for example, this included how to obtain GNU Emacs and Richard Stallman's political essay "The GNU Manifesto". The XEmacs manuals, which were inherited from older GNU Emacs manuals when the fork occurred, have the same license. The newer versions of the GNU Emacs documentation, meanwhile, uses the GNU Free Documentation License and makes use of "invariant sections" to require the inclusion of the same documents, additionally requiring that the manuals proclaim themselves as *GNU Manuals*.

Using Emacs

Commands

From the Unix shell, a file can be opened for editing by typing "emacs [filename]". If the filename you entered does not exist a file will be created with that name. For example "emacs xorg.conf" will edit the xorg.conf file in the current directory, if it exists. However, Emacs documentation recommends starting Emacs without a file name, to avoid the bad habit of starting a separate Emacs for each file you edit. Visiting all files in a single Emacs process is the way to get the full benefit of Emacs.

In the normal editing mode, Emacs behaves just like other text editors: the character keys (*a*, *b*, *c*, *1*, *2*, *3*, etc.) insert

the corresponding characters, the arrow keys move the editing point, backspace deletes text, and so forth. Other commands are invoked with modified keystrokes, pressing the control key and/or the meta key/alt key in conjunction with a regular key. Every editing command is actually a call to a function in the Emacs Lisp environment. Even a command as simple as typing a to insert the character *a* involves calling a function--in this case, `self-insert-command`.

Some of the basic commands are shown below. More can be found at List of Emacs commands. The control key [Ctrl] is denoted by a capital *C*, and the meta or alt [Alt] key by a capital *M*.

Command	Keystroke	Description
forward-word	M-f	Move forward past one word.
search-word	C-s	Search a word in the buffer.
undo	C-/	Undo last change, and prior changes if pressed repeatedly.
keyboard-quit	C-g	Abort the current command.
fill-paragraph	M-q	Wrap text in ("fill") a paragraph.
find-file	C-x C-f	Visit a file (you specify the name) in its own editor buffer.
save-buffer	C-x C-s	Save the current editor buffer in its visited file.
write-file	C-x C-w	Save the current editor buffer as a file with the name you specify.
save-buffers-kill-emacs	C-x C-c	Offer to save changes, then exit Emacs.
set-marker	C-[space]/C-@	Set a marker from where you want to cut or copy.
cut	C-w	Cut all text between the marker and the cursor.
copy	M-w	Copy all text between the marker and the cursor.
paste	C-y	Paste text from the emacs clipboard
kill buffer	C-x k	Kill the current buffer

Alternatively, if a user would prefer IBM Common User Access style keys, "cua-mode" can be used. This has been a third-party package up to, and including, GNU Emacs 21, but is included in GNU Emacs 22 (beta).

Note that the commands `save-buffer` and `save-buffers-kill-emacs` use *multiple* modified keystrokes. For example, `C-x C-c` means: while holding down the control key, press *x*; then, while holding down the control key, press *c*. This technique, allowing more commands to be bound to the keyboard than with the use of single keystrokes alone, was popularized by Emacs, which got it from TECMAC, one of the TECO macro collections that immediately preceded Emacs. It has since made its way into modern code editors like Visual Studio.

When Emacs is running a graphical interface, many commands can be invoked from the menubar or toolbar instead of using the keyboard. However, many experienced Emacs users prefer to use the keyboard because it is faster and more convenient once the necessary keystrokes have been memorized.

Some Emacs commands work by invoking an external program (such as `ispell` for spell-checking or `gcc` for program compilation), parsing the program's output, and displaying the result in Emacs.

Minibuffer

The *minibuffer*, normally the bottommost line, is where Emacs requests information. Text to target in a search, the

name of a file to read or save and similar information is entered in the minibuffer. When applicable, tab completion is usually available.

File management and display

Emacs keeps text in objects called *buffers*. The user can create new buffers and dismiss unwanted ones, and several buffers can exist at the same time. Most buffers contain text loaded from text files, which the user can edit and save back to disk. Buffers are also used to store temporary text, such as the documentation strings displayed by the *help* library.

In both text terminal and graphical modes, Emacs is able to split the editing area into separate sections (referred to since 1975 as "windows", which can be confusing on systems that have another concept of "windows" as well), so that more than one buffer can be displayed at a time. This has many uses. For example, one section can be used to display the source code of a program, while another displays the results from compiling the program. In graphical environments, Emacs can also launch multiple graphical-environment windows, known as "frames" in the context of Emacs.

Emacs Pinky

Because of Emacs' dependence on the modifier keys, in particular the control key is pressed with the little finger ("pinky"), heavy Emacs users have experienced pain in their pinky fingers (see repetitive strain injury and fat-finger). This has been dubbed the "Emacs Pinky", and vi advocates often cite it as a reason to switch to vi. To alleviate this situation, many Emacs users transpose the left control key and the left caps-lock key or define both as control keys. There are also Kinesis's Contoured Keyboard available which reduce the strain by moving the modifier keys altogether so that they are in a position to be easily pushed by the thumb, and Microsoft Natural keyboard that has large modifier keys placed symmetrically on both sides of the keyboard so that they can be pressed with the palm.

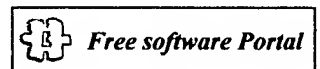
Distractions

In addition to its many features, GNU Emacs includes a variety of unusual distractions designed to amuse and/or annoy.

- M-x *life* renders Conway's Game of Life
- M-x *gomoku* launches a game of Gomoku
- M-x *psychoanalyze-pinhead* pipes Zippy the Pinhead quotes through ELIZA Doctor
- M-x *spook* outputs a string of random words designed to distract anyone from the NSA who might be listening in

See also

- Comparison of text editors
- GNU TeXmacs
- List of text editors
- List of Unix programs



References

1. ^ Thomas Widmann (2005-06-13). "'Emacs' defined in Collins English Dictionary". comp.emacs. (Google

Groups). Retrieved on 2006-09-27.

2. ^ Stephen J., Turnbull. XEmacs vs. GNU Emacs. Retrieved on 2006-09-27.
3. ^ <http://www.stifflog.com/2006/10/16/stiff-asks-great-programmers-answer/>
4. ^ <http://www.meadowy.org/meadow/pukiwiki-en/>
5. ^ Carbon Emacs Package. Retrieved on 2006-09-27.
6. ^ Aquamacs is an easy-to-use, Mac-style Emacs for Mac OS X. Retrieved on 2006-09-27.
7. ^ B, Ramprasad (2005-06-24). GNU Emacs FAQ For Windows 95/98/ME/NT/XP and 2000. Retrieved on 2006-09-27.
8. ^ Borgman, Lennart (2006). EmacsW32 Home Page. Retrieved on 2006-09-27.
9. ^ GNU Emacs on Windows. Franz Inc. (2006). Retrieved on 2006-09-27.

Bibliography

- Ciccarelli, Eugene (1978). *An Introduction to the Emacs Editor*. Cambridge, Massachusetts: MIT Artificial Intelligence Laboratory. AIM-447. PDF HTML
- Stallman, Richard M. (1979, updated 1981). *EMACS: The Extensible, Customizable, Self-Documenting Display Editor*. Cambridge Massachusetts: MIT Artificial Intelligence Laboratory. AIM-519A. PDF HTML
- Stallman, Richard M (2002). *GNU Emacs Manual*, 15th edition, GNU Press. ISBN 1-882114-85-X.
- Stallman, Richard M (2002). My Lisp Experiences and the Development of GNU Emacs. Retrieved on 2007-02-01.
- Chassel, Robert J. (2004). *An Introduction to Programming in Emacs Lisp*. GNU Press. ISBN 1-882114-56-6.
- Glickstein, Bob (1997 (April)). *Writing GNU Emacs Extensions*. O'Reilly & Associates. 1-56592-261-1.
- Cameron, Debra; Elliott, James; Loy, Marc; Raymond, Eric; Rosenblatt, Bill (2004 (December)). *"Learning GNU Emacs, 3rd Edition"*. O'Reilly & Associates. ISBN 0-596-00648-9.
- Greenberg, Bernard S. (1979). *"Multics Emacs: The History, Design and Implementation"*.
- Finseth, Craig A. (1991). *The Craft of Text Editing -or- Emacs for the Modern World*. Springer-Verlag & Co. ISBN 978-1-4116-8297-9.
- Zawinski, Jamie (1999, updated 2005-06-21). Emacs Timeline. Retrieved on 2006-09-30.

External links

- The GNU Emacs homepage
- List of Emacs implementations
- EmacsWiki – community site dedicated to documenting and discussing Emacs
- Graphical tutorial of Emacs in pdf format
- Printable Key Bindings Reference Card Plain Text

GNU Project

[hide]

History: GNU Manifesto • GNU Project • Free Software Foundation (FSF) • History of free software

GNU licenses: GNU General Public License (GPL) • GNU Lesser General Public License (LGPL) • GNU Free Documentation License (FDL)

Software: GNU operating system • bash • GNU Compiler Collection • **Emacs** • GNU C Library • Coreutils • GNU build system • *other GNU packages and programs*

Speakers: Robert J. Chassel • Loïc Dachary • Ricardo Galli • Georg C. F. Greve • Federico Heinz • Bradley M. Kuhn • Eben Moglen • Richard Stallman • Len Tower

Retrieved from "<http://en.wikipedia.org/wiki/Emacs>"

Categories: Articles with unsourced statements since February 2007 | All articles with unsourced statements | GNU

[project](#) | [Emacs](#) | [Free file comparison tools](#) | [Free integrated development environments](#) | [Free text editors](#) | [GNU project software](#) | [Integrated development environments](#) | [Linux text editors](#) | [Mac OS text editors](#) | [Mac OS X text editors](#) | [OpenVMS text editors](#) | [Windows text editors](#) | [Linux integrated development environments](#)

- This page was last modified 15:12, 22 April 2007.
- All text is available under the terms of the GNU Free Documentation License. (See **Copyrights** for details.) Wikipedia® is a registered trademark of the Wikimedia Foundation, Inc., a US-registered 501(c)(3) tax-deductible nonprofit charity.